arXiv:1609.01872v1 [stat.ML] 7 Sep 2016

Chaining Bounds for Empirical Risk Minimization

Gábor Balázs and Csaba Szepesvári
Department of Computing Science, University of Alberta,
Edmonton, Alberta, T6G 2E8, Canada

András György
Department of Electrical and Electronic Engineering,
Imperial College London, London, SW7 2BT, United Kingdom

September 8, 2016

Abstract

This paper extends the standard chaining technique to prove excess risk upper bounds for empirical risk minimization with random design settings even if the magnitude of the noise and the estimates is unbounded. The bound applies to many loss functions besides the squared loss, and scales only with the sub-Gaussian or subexponential parameters without further statistical assumptions such as the bounded kurtosis condition over the hypothesis class. A detailed analysis is provided for slope constrained and penalized linear least squares regression with a sub-Gaussian setting, which often proves tight sample complexity bounds up to logartihmic factors.

1 Introduction

This paper extends the standard chaining technique (e.g., Pollard, 1990; Dudley, 1999; Györfi et al., 2002; Boucheron et al., 2012) to prove high-probability excess risk upper bounds for empirical risk minimization (ERM) for random design settings even if the magnitude of the noise and the estimates is unbounded. Our result (Theorem 1) covers bounded settings (Bartlett et al., 2005; Koltchinskii, 2011), extends to sub-Gaussian or even subexponential noise (van de Geer, 2000; Györfi and Wegkamp, 2008), and handles hypothesis classes with unbounded magnitude (Lecué and Mendelson, 2013; Mendelson, 2014; Liang et al., 2015). Furthermore, it applies to many loss functions besides the squared loss, and does not need additional statistical assumptions such as the bounded kurtosis of the transformed covariates over the hypothesis class, which prevent the latest developments to provide tight excess risk bounds for many sub-Gaussian cases (Section 1.2).

To demonstrate the effectiveness of our method for such unbounded settings, we use our general excess risk bound (Theorem 1) to provide a detailed analysis for linear least squares estimators using quadratic slope constraint and penalty with sub-Gaussian noise and domain for the random design, nonrealizable setting (Section 3). Our result for the slope constrained case extends Theorem A of Lecué and Mendelson (2013) and nearly proves the conjecture of Shamir (2015), while our treatment for the penalized case (ridge regression) is comparable to the work of Hsu et al. (2014).

The rest of this section introduces our notation through the formal definition of the regression problem and ERM estimators (Section 1.1), and discusses the limitations of

current excess risk upper bounds in the literature (Section 1.2). Then, we provide our main result in Section 2 to upper bound the excess risk of ERM estimators, and discuss its properties for various settings including many loss functions besides the squared loss. Next, Section 3 provides a detailed analysis for linear least squares estimators including the slope constrained case (Section 3.1) and ridge regression (Section 3.2). Finally, Section 4 proves our main result (Theorem 1).

1.1 Empirical risk minimization

For the formal definition of a regression problem, consider a probability distribution μ over some set $\mathbb{X} \times \mathbb{R}$ with some *domain* $\mathbb{X}$ being a separable Hilbert space,[1] a *loss function* $\ell : \mathbb{R} \times \mathbb{R} \to [0, \infty)$, and a *reference class* $\mathcal{F}_* \subseteq \{\mathbb{X} \to \mathbb{R}\} \doteq \{f \,|\, f : \mathbb{X} \to \mathbb{R}\}$.

The task of a regression estimator is to produce a function $f \in \{\mathbb{X} \to \mathbb{R}\}$ based on a *training sample* $\mathcal{D}_n \doteq \{(\boldsymbol{\mathcal{X}}_1, \mathcal{Y}_1), \ldots, (\boldsymbol{\mathcal{X}}_n, \mathcal{Y}_n)\}$ of $n \in \mathbb{N}$ pairs $(\boldsymbol{\mathcal{X}}_i, \mathcal{Y}_i) \in \mathbb{X} \times \mathbb{R}$ independently sampled from μ (in short $\mathcal{D}_n \sim \mu^n$), such that the prediction error, $\ell(\mathcal{Y}, f(\boldsymbol{\mathcal{X}}))$, is "small" on a new instance $(\boldsymbol{\mathcal{X}}, \mathcal{Y}) \sim \mu$ with respect to ℓ.

The *risk* of function $f \in \{\mathbb{X} \to \mathbb{R}\}$ is defined as $R_\mu(f) \doteq \mathbb{E}[\ell(\mathcal{Y}, f(\boldsymbol{\mathcal{X}}))]$ and the cost of using a fixed function f is measured by the *excess risk* with respect to $\mathcal{F}_*$:

$$L_\mu(f, \mathcal{F}_*) \doteq R_\mu(f) - \min_{g \in \mathcal{F}_*} R_\mu(g)\,.$$

We also use the notation $L_\mu(f, g) \doteq R_\mu(f) - R_\mu(g)$ for any $f, g \in \{\mathbb{X} \to \mathbb{R}\}$, hence we can write $L_\mu(f, \mathcal{F}_*) = L_\mu(f, f_*)$ for any $f_* \in \operatorname{argmin}_{g \in \mathcal{F}_*} R_\mu(g)$.[2]

An *estimator* h_n is a sequence of functions $h \doteq (h_n)_{n \in \mathbb{N}}$, where $h_n : (\mathbb{X} \times \mathbb{R})^n \to \{\mathbb{X} \to \mathbb{R}\}$ maps the data $\mathcal{D}_n$ to an estimate $f_n \doteq h_n(\mathcal{D}_n)$. These estimates lie within some *hypothesis class* $\mathcal{F}_n \subseteq \{\mathbb{X} \to \mathbb{R}\}$, that is $f_n \in \mathcal{F}_n$, where $\mathcal{F}_n$ might depend on the random sample $\mathcal{D}_n$.

Then, for a regression problem specified by $(\ell, \mu, \mathcal{F}_*)$, the goal of an estimator h_n is to produce estimates which minimize the excess risk $L_\mu(f_n, \mathcal{F}_*)$ with high-probability or in expectation, where the random event is induced by the random sample $\mathcal{D}_n$ and the possible randomness of the estimator h_n.

In this paper, we consider ERM estimators. Formally, f_n is called an *α-approximate β-penalized ERM estimate* with respect to the class $\mathcal{F}_n$, in short $f_n \in (\alpha, \beta)\text{-ERM}(\mathcal{F}_n)$, when $f_n \in \mathcal{F}_n$ and

$$R_n(f_n) + \beta(f_n) \le \inf_{f \in \mathcal{F}_n} R_n(f) + \beta(f) + \alpha\,, \tag{1}$$

where $R_n(f) \doteq \frac{1}{n}\sum_{i=1}^n \ell(\mathcal{Y}_i, f(\boldsymbol{\mathcal{X}}_i))$ is the *empirical risk* of function $f \in \{\mathbb{X} \to \mathbb{R}\}$, $\beta : \mathcal{F}_n \to [0, \infty)$ is a *penalty function* and $\alpha \ge 0$ is an *error term*. All α, β, and $\mathcal{F}_n$ might depend on the sample $\mathcal{D}_n$. When the penalty function is zero (that is $\beta \equiv 0$), we simply write $f_n \in \alpha\text{-ERM}(\mathcal{F}_n)$. If both $\alpha = 0$ and $\beta \equiv 0$, we say $f_n \in \text{ERM}(\mathcal{F}_n)$.

1.2 Limitations of current methods

Now we provide a simple regression problem class for which we are not aware of any technique in the literature that could provide a tight excess risk bound up to logarithmic factors

[1] All sets and functions considered are assumed to be measurable as necessary. To simplify the presentation, we omit these conditions by noting here that all the measurability issues can be overcome using standard techniques as we work with separable Hilbert spaces (e.g., Dudley, 1999, Chapter 5).

[2] A straightforward limiting argument can be used if the minimums are not attained for f_*.

for empirical risk minimization. Consider the following problem set:

$$\mathbb{M}_{\text{bg}} \doteq \big\{\mu \,\big|\, (\boldsymbol{\mathcal{X}}, \mathcal{Y}) \sim \mu,\ \boldsymbol{\mathcal{X}} = [\mathcal{W}\ \mathcal{Z}\ 1]^\top \in \mathbb{X} \subset \mathbb{R}^3,\\ \mathcal{W} \sim \mathcal{N}(0,1),\ \mathbb{P}\{\mathcal{Z} = -p\} = 1-p,\ \mathbb{P}\{\mathcal{Z} = 1-p\} = p,\ p \in (0,1),\\ \mathcal{Y} = f_*(\boldsymbol{\mathcal{X}}) + 1/2 + \xi \text{ a.s.},\ f_*(\boldsymbol{x}) = [0\ 0\ 1/2]\boldsymbol{x},\ \boldsymbol{x} \in \mathbb{X},\ \xi \sim \mathcal{N}(0,1)\big\},$$

where $\mathbb{X} \doteq \mathbb{R} \times [-1,1] \times \{1\}$, a.s. stands for almost surely, and $\mathcal{N}(0,1)$ denotes the centered Gaussian distribution with unit variance. Furthermore, define the linear function class $\mathcal{F} = \{f \,|\, f(\boldsymbol{x}) = \boldsymbol{a}^\top \boldsymbol{x},\ \|\boldsymbol{a}\| \le 1/2\}$, and consider the squared loss $\ell = \ell_{\text{sq}}$ defined as $\ell_{\text{sq}}(y, \hat{y}) \doteq |y - \hat{y}|^2$ for all $y, \hat{y} \in \mathbb{R}$. Notice that $f_* \in \mathcal{F}$ and $f_* = \operatorname{argmin}_{f \in \mathcal{F}} R_\mu(f)$ for all $\mu \in \mathbb{M}_{\text{bg}}$. Then, we discuss various techniques from the literature which aim to bound the "performance" of an estimate $f_n \in \text{ERM}(\mathcal{F})$.

Because here we have a random design setting, the results of van de Geer (2000, Theorems 9.1 and 9.2) do not apply. Moreover, the regression function cannot be represented by the class $\mathcal{F}$, so the methods of Györfi et al. (2002, Theorem 11.3), and Györfi and Wegkamp (2008, Corollary 1) do not provide an excess risk bound.

As the domain $\mathbb{X}$ is unbounded, so does the range of any nonzero function in $\mathcal{F}$. Additionally, the squared loss ℓ_{sq} is neither Lipschitz, nor bounded on the range of response $\mathcal{Y}$ which is the whole real line $\mathbb{R}$. Hence, the techniques including Bartlett et al. (2005, Corollary 5.3), Koltchinskii (2011, Theorem 5.1), Mehta and Williamson (2014, Theorem 6), Grünwald and Mehta (2016, Theorem 14 with Proposition 4) fail to provide any rate for this case.

We also mention the work of van der Vaart and Wellner (2011, Theorem 3.2), which, although works for this setting, can only provide an $O(\sqrt{\ln(n)/n})$ rate for sample size n, which can be improved to $O(\ln(n)/n)$ by our result (Theorem 1 and Lemma 3).

Next, denote the *kurtosis about the origin* by $\mathbb{K}_0[\mathcal{W}] \doteq \mathbb{E}[\mathcal{W}^4]/\mathbb{E}[\mathcal{W}^2]^2$ for some random variable $\mathcal{W}$, and consider the recent developments of Lecué and Mendelson (2013, Theorem A), and Liang et al. (2015, Theorem 7). These results need that the kurtosis of the random variables $\Delta_{f,f_*} \doteq f(\boldsymbol{\mathcal{X}}) - f_*(\boldsymbol{\mathcal{X}})$ is bounded for any $f \in \mathcal{F} \setminus \{f_*\}$. However, observe that $\mathbb{K}_0[\Delta_{f,f_*}] \ge (1-p)^2/p$ for any function $f(\boldsymbol{x}) = [0\ a\ 0]\boldsymbol{x}$ with $0 < |a| \le 1/2$, which can be arbitrarily large as p gets close to zero.

Finally, we mention the result of Mendelson (2014, Theorem 2.2), which bounds the squared deviation of the ERM estimator f_n and f_*, that is $\mathbb{E}\big[|f_n(\boldsymbol{\mathcal{X}}) - f_*(\boldsymbol{\mathcal{X}})|^2\big]$. However, as pointed out by Shamir (2015, Section 1), this can be arbitrarily smaller than the excess risk $L_\mu(f_n, f_*)$ for functions $f_a(\boldsymbol{x}) = [0\ 0\ a]\boldsymbol{x}$ with $|a| \le 1/2$, which means that $L_\mu(f_a, f_*)/\mathbb{E}\big[|f_a(\boldsymbol{\mathcal{X}}) - f_*(\boldsymbol{\mathcal{X}})|^2\big] \to \infty$ as $a \to 1/2$.

1.3 Highlights of our technique

Our excess risk bound builds on the development of inexact oracle inequalities for ERM estimators (e.g., Györfi et al., 2002, Theorem 11.5), which uses the decomposition

$$\begin{aligned} L_\mu(f_n, f_*) &= L_\mu(f_n, f_*) - cL_n(f_n, f_*) + cL_n(f_n, f_*) \\ &\le \sup_{f \in \mathcal{F}_n} \big\{L_\mu(f, f_*) - cL_n(f, f_*)\big\} + c \inf_{f \in \mathcal{F}_n} L_n(f, f_*)\,, \end{aligned}$$

with some $c > 1$. Then the random variables $L_\mu(f, f_*) - cL_n(f, f_*)$ for all $f \ne f_*$, having a negative bias, often satisfy a moment condition (C_5), which we cannot guarantee for $c = 1$. Using this moment condition, we can augment the chaining technique (e.g., Pollard, 1990, Section 3) with an extra initial step, which provides a new $O(1/n)$ term in the bound. This

new term can be balanced with the (truncated) entropy integral, so tightening the bound significantly in many cases.

By defining f_* as a reference function (instead of regression function), the inexact oracle inequalities become exact when $f_* \in \mathcal{F}_n$. In fact, the notion of exact and inexact becomes meaningless as long as the approximation error between $\mathcal{F}_n$ and f_* is kept under control and incorporated into the bound as it is often done for sieved estimators (e.g., van de Geer, 2000, Section 10.3).

To prove the moment condition (C_5) for a reference function f_* and a hypothesis class $\mathcal{F}_n$, we use Bernstein's inequality (Lemma 14) with the Bernstein condition (3). These tools are standard, however we have to use Bernstein's inequality for the sub-Gaussian random variable Δ_{f,f_*} so that $\mathbb{E}[\Delta_{f,f_*}^2]$ appears in the bound. A naive way to do this would require the kurtosis $\mathbb{K}_0[\Delta_{f,f_*}]$ to be bounded for all $f \in \mathcal{F}_n$, which cannot be guaranteed in many cases (Section 1.2). Hence, we use a truncation technique (Lemma 15) that pushes the kurtosis bound $\sup_{f\in\mathcal{F}_n} \mathbb{K}_0[\Delta_{f,f_*}]$ under a logarithmic transformation, which can be eliminated by considering functions $f \in \mathcal{F}_n$ with excess risk $L_\mu(f, f_*)$ bounded away from zero.

The Bernstein condition (3) has been well-studied for strongly-convex loss functions (e.g., Bartlett et al., 2006, Lemma 7), by exploiting that strong-convexity provides an upper bound to the quadratic function. However, because it is enough for our technique to consider functions with excess risk bouded away from zero, we can use the Bernstein condition for any Lipschitz loss function (Section 2.2.1) by scaling its parameters depending on the sample size n and balancing the appropriate terms in the excess risk bound (Theorem 1). In many cases, this provides an alternative way for deriving excess risk bounds for other loss functions without using the entropy integral.

2 Excess risk upper bound

Here we are going to state our excess risk upper bound for ERM estimators.

Our result requires a few conditions to be satisfied by the random variables $\mathcal{Z}(f,g) \doteq \ell(\mathcal{Y}, f(\boldsymbol{\mathcal{X}})) - \ell(\mathcal{Y}, g(\boldsymbol{\mathcal{X}}))$ with $f, g \in \{\mathbb{X} \to \mathbb{R}\}$, which are related to the excess risk through $L_\mu(f,g) = \mathbb{E}[\mathcal{Z}(f,g)]$. Similarly, we use the *empirical excess risk* defined as $L_n(f,g) \doteq R_n(f) - R_n(g) = \frac{1}{n}\sum_{i=1}^n \mathcal{Z}_i(f,g)$, where $\mathcal{Z}_i(f,g) \doteq \ell(\mathcal{Y}_i, f(\boldsymbol{\mathcal{X}}_i)) - \ell(\mathcal{Y}_i, g(\boldsymbol{\mathcal{X}}_i))$.

We also use subexponential random variables ($d = 1$) and vectors $\boldsymbol{\mathcal{W}} \in \mathbb{R}^d$ characterized by the Ψ_q-Orlicz norm with $q \geq 1$ defined as $\|\boldsymbol{\mathcal{W}}\|_{\Psi_q} \doteq \inf\{B > 0 : \mathbb{E}[\Psi_q(\boldsymbol{\mathcal{W}}/B)] \leq 1\}$, where $\Psi_q(\boldsymbol{x}) \doteq e^{\|\boldsymbol{x}\|^q} - 1$, $\boldsymbol{x} \in \mathbb{R}^d$, $\|\cdot\|$ is the Euclidean norm, and $\inf \emptyset = \infty$. The properties of random vectors $\boldsymbol{\mathcal{W}}$ with $\|\boldsymbol{\mathcal{W}}\|_{\Psi_q} < \infty$ are reviewed in Appendix A.

Furthermore, we need covering numbers and entropies. Let $(\mathcal{F}, \psi)$ be a nonempty metric space and $\epsilon \geq 0$. The set $\{f_1, \ldots, f_k\} \subseteq \mathcal{F}$ is called an (internal) *ϵ-cover* of $\mathcal{F}$ under ψ if the ψ-balls of centers $\{f_1, \ldots, f_k\}$ and radius ϵ cover $\mathcal{F}$: for any $f \in \mathcal{F}$, $\min_{i=1,\ldots,k} \psi(f, f_i) \leq \epsilon$. The *$\epsilon$-covering number* of $\mathcal{F}$ under ψ, denoted by $\mathcal{N}_\psi(\epsilon, \mathcal{F})$, is the cardinality of the ϵ-cover with the fewest elements:

$$\mathcal{N}_\psi(\epsilon, \mathcal{F}) \doteq \inf\left\{k \in \mathbb{N} \,\middle|\, \exists f_1, \ldots, f_k \in \mathcal{F} : \sup_{f\in\mathcal{F}} \min_{i=1,\ldots,k} \psi(f, f_i) \leq \epsilon\right\}$$

with $\inf \emptyset \doteq \infty$. Further, the *$\epsilon$-entropy* of $\mathcal{F}$ under ψ is defined as the logarithm of the covering number, $\mathcal{H}_\psi(\epsilon, \mathcal{F}) \doteq \ln \mathcal{N}_\psi(\epsilon, \mathcal{F})$.

Finally, our upper bound on the excess risk of ERM estimates is the following:

Theorem 1. *Consider a regression problem $(\ell, \mu, \mathcal{F}_*)$ with an i.i.d. training sample $\mathcal{D}_n \sim \mu^n$. Let $\mathcal{F}_n \subseteq \{\mathbb{X} \to \mathbb{R}\}$ be a hypothesis class which might depend on the data $\mathcal{D}_n$, and let $f_n \in (\alpha, \beta)$-ERM($\mathcal{F}_n$). Further, let $\hat{\mathcal{F}}_n, \mathcal{F} \subseteq \{\mathbb{X} \to \mathbb{R}\}$ be two function classes, where $\hat{\mathcal{F}}_n$ might depend on $\mathcal{D}_n$, but $\mathcal{F}$ might depend on the sample only through its size n. Finally, suppose that the following conditions hold for some metric $\psi : \mathcal{F} \times \mathcal{F} \to [0, \infty)$, $\gamma \in (0, 1)$, $f_* \in \operatorname{argmin}_{f \in \mathcal{F}_*} R_\mu(f)$, and $\mathcal{F}(r_0) \doteq \{f \in \mathcal{F} : L_\mu(f, \mathcal{F}_*) > r_0/n\}$ for some $r_0 \geq 0$:*

(C_1) the enclosement $\hat{\mathcal{F}}_n \subseteq \mathcal{F}_n \subseteq \mathcal{F}$ holds with probability at least $1 - \gamma/4$,

(C_2) there exists $B_* \geq 0$ such that $\mathbb{P}\big\{ \inf_{f \in \hat{\mathcal{F}}_n} L_n(f, f_*) + \beta(f) + \alpha \leq B_* \big\} \geq 1 - \gamma/4$,

(C_3) there exists $G : \mathbb{X} \times \mathbb{R} \to [0, \infty)$ and $T \geq 0$ such that $\mathcal{Z}(f, g) \leq G(\boldsymbol{\mathcal{X}}, \mathcal{Y})\, \psi(f, g)$ a.s. for all $f, g \in \mathcal{F}(r_0)$, and $\mathbb{P}\big\{ \mathbb{E}[G(\boldsymbol{\mathcal{X}}, \mathcal{Y})] + \frac{1}{n} \sum_{i=1}^n G(\boldsymbol{\mathcal{X}}_i, \mathcal{Y}_i) \leq T \big\} \geq 1 - \gamma/4$,

(C_4) there exists $S \in (0, \infty]$ and $q \in \{1, 2\}$ such that $\| \mathbb{E}[\mathcal{Z}(f, g)] - \mathcal{Z}(f, g) \|_{\Psi_q} \leq S\, \psi(f, g)$ holds for all $f, g \in \mathcal{F}(r_0)$,

(C_5) there exist $r \in (0, 1]$ and $\theta > 0$ such that

$$\sup_{f \in \mathcal{F}(r_0)} \mathbb{E}\big[e^{(r/\theta)\mathbb{E}[\mathcal{Z}(f, f_*)] - (1/\theta)\mathcal{Z}(f, f_*)} \big] \leq 1 \,.$$

Then for all $\epsilon \geq \delta \geq 0$, we have with probability at least $1 - \gamma$ that

$$L_\mu(f_n, \mathcal{F}_*) \leq \frac{1}{r} \left(\theta \frac{\mathcal{H}_\psi\big(\epsilon, \mathcal{F}(r_0)\big) + \ln(\frac{4}{\gamma})}{n} + \frac{32 S}{\sqrt{n}} \int_\delta^\epsilon \sqrt[q]{2 \mathcal{H}_\psi\big(z, \mathcal{F}(r_0)\big) + \ln\Big(\frac{32\epsilon}{z\gamma}\Big)}\, dz + 8\delta T + B_* \right) + \frac{r_0}{n} .$$

Furthermore, the result holds without (C_4), that is using $S = \infty$, $\epsilon = \delta$, and $\infty \cdot 0 = 0$.

The proof of Theorem 1 is presented in Section 4.

We point out that (C_1) disappears when one sets $\hat{\mathcal{F}}_n = \mathcal{F}_n = \mathcal{F}$ as it is usually done in the literature. However, this is an implicit assumption that either $\mathbb{E}\boldsymbol{\mathcal{X}}$ is small enough to be negligible (i.e., $\mathbb{E}\boldsymbol{\mathcal{X}} \approx \mathbf{0}$), or equivalently the estimator knows the value of $\mathbb{E}\boldsymbol{\mathcal{X}}$. By choosing the sets $\hat{\mathcal{F}}_n, \mathcal{F}_n, \mathcal{F}$ to be slightly different, Theorem 1 covers the practical case when an estimator approximates $\mathbb{E}\boldsymbol{\mathcal{X}}$ and $\mathbb{E}\mathcal{Y}$ by their empirical versions $\overline{\boldsymbol{\mathcal{X}}} \doteq \frac{1}{n} \sum_{i=1}^n \boldsymbol{\mathcal{X}}_i$ and $\overline{\mathcal{Y}} \doteq \frac{1}{n} \sum_{i=1}^n \mathcal{Y}_i$, respectively, so uses a data-dependent hypothesis class $\mathcal{F}_n$.

Notice that if $f_* \in \hat{\mathcal{F}}_n$, then ($C_2$) reduces to bounding the penalty and error terms, that is proving $\inf_{f \in \hat{\mathcal{F}}_n} \beta(f) + \alpha \leq B_*$ with probability at least $1 - \gamma/4$. When $f_n \in \mathrm{ERM}(\mathcal{F}_n)$, which is a usual setting in the literature, (C_2) is immediately satisfied by $B_* = 0$. In this case Theorem 1 is an *exact oracle inequality* (e.g., Lecué and Mendelson, 2013, Eq. 1.1).

Furthermore, observe that Theorem 1 uses metric ψ for the entropy $\mathcal{H}_\psi(\cdot, \mathcal{F}(r_0))$, which is related to the loss function ℓ through (C_3) and (C_4). This allows us to apply the result to estimates f_n with unbounded magnitude, which can be parametrized by some bounded space. In such case ψ is defined on the bounded parameter space which keeps the entropy finite.

In the following sections we provide a detailed analysis for the *moment condition* (C_5), showing that it holds for many practical settings and loss functions besides the squared loss. We note that (C_5) is very similar to the *stochastic mixability condition* of Mehta and Williamson (2014, Section 2.1), which is equivalent to (C_5) with $r = 0$ and $r_0 = 0$.

Finally, we mention that if the conditions of Theorem 1 hold for all $\gamma \in (0, 1)$, we can transform the result to an expected excess risk bound. To see this, suppose that

$\mathbb{P}\{L_\mu(f_n, \mathcal{F}_*) \le b + \frac{c \ln^m(1/\gamma)}{n}\} \ge 1 - \gamma$ holds for all $\gamma \in (0,1)$, some $b, c > 0$, and some $m \in \mathbb{N}$. Then setting $\gamma = e^{-(n\,t/c)^{1/m}}$ for any $t > 0$, we get

$$\begin{aligned}\mathbb{E}[L_\mu(f_n, \mathcal{F}_*)] - b &\le \mathbb{E}\big[\max\{0, L_\mu(f_n, \mathcal{F}_*) - b\}\big] \\ &= \int_0^\infty \mathbb{P}\{L_\mu(f_n, \mathcal{F}_*) > b + t\}\, dt \le \int_0^\infty e^{-(n\,t/c)^{1/m}}\, dt = \frac{m!\, c}{n}\,, \end{aligned} \tag{2}$$

where the expectation is taken with respect to the random sample $\mathcal{D}_n$ and the potential extra randomness of the estimator h_n producing $f_n = h_n(\mathcal{D}_n)$.

2.1 Bounded losses

We start with a simple case when the loss function ℓ is bounded, which implies that $|\mathcal{Z}(f, f_*)| \le B$ holds for some $B > 0$. Now notice that the random variable in the exponent of (C_5), that is $r\mathbb{E}[\mathcal{Z}(f, f_*)] - \mathcal{Z}(f, f_*)$, has a negative expected value $(r-1)\mathbb{E}[\mathcal{Z}] < 0$. Furthermore, $\mathbb{E}[\mathcal{Z}(f, f_*)]$ is bounded away from zero as $\mathbb{E}[\mathcal{Z}(f, f_*)] > r_0/n$ by the definition of $\mathcal{F}(r_0)$. Then, combining these observations with Hoeffding's lemma, we get the following result:

Lemma 2. *Suppose that* $\sup_{f \in \mathcal{F}(r_0)} |\mathcal{Z}(f, f_*)| \le B$ *a.s. holds for some* $B > 0$. *Then* $(\mu, \ell, \mathcal{F}(r_0), f_*)$ *satisfies* (C_5) *with any* $r \in (0,1)$ *and* $\theta \ge nB^2/(2(1-r)r_0)$.

Proof. Fix any $f \in \mathcal{F}(r_0)$, and set $\mathcal{Z}_f \doteq \mathcal{Z}(f, f_*)$. Then, apply Hoeffding's lemma to the bounded random variable $\mathcal{Z}_f$ to get

$$\mathbb{E}\big[e^{(r/\theta)\mathbb{E}[\mathcal{Z}_f] - (1/\theta)\mathcal{Z}_f}\big] = e^{(r-1)\mathbb{E}[\mathcal{Z}_f]/\theta}\, \mathbb{E}\big[e^{(\mathbb{E}[\mathcal{Z}_f] - \mathcal{Z}_f)/\theta}\big] \le e^{(r-1)\mathbb{E}[\mathcal{Z}_f]/\theta + B^2/(2\theta^2)} \le 1\,,$$

as $\frac{(r-1)\mathbb{E}[\mathcal{Z}_f]}{\theta} + \frac{B^2}{2\theta^2} < \mathbb{E}[\mathcal{Z}_f]\big(\frac{(r-1)}{\theta} + \frac{n}{2r_0\theta^2}\big) \le 0$ by $1 < \frac{n}{r_0}\mathbb{E}[\mathcal{Z}_f]$ due to the definiton of $\mathcal{F}(r_0)$ and $\mathbb{E}[\mathcal{Z}_f] = L_\mu(f, f_*)$. ∎

However, as θ scales with n/r_0, one should choose $r_0 > 0$ to balance the appropriate terms of Theorem 1. This can be achieved by setting $r_0 = \Theta\big(B\sqrt{n\,\mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0))}\big)$ which balances $\frac{\theta}{n}\mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0)) = \Theta\big(\frac{B^2}{r_0}\mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0))\big)$ with r_0/n. This way the bound of Theorem 1 scales with $B\sqrt{\mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0))/n}$, which cannot be improved in general (e.g., Bartlett et al., 2005, Section 5).

2.2 Unbounded losses

We show that the moment condition (C_5) is often implied by the *Bernstein condition* (e.g., Lecué and Mendelson, 2013, Definition 1.2), which is said to be satisfied by $(\mu, \ell, \mathcal{F}, f_*)$ if there exists $C > 0$ such that for all $f \in \mathcal{F}$, we have

$$\mathbb{E}\big[\Delta_{f,f_*}^2\big] \le C\,\mathbb{E}\big[\mathcal{Z}(f, f_*)\big]\,, \quad \Delta_{f,f_*} \doteq f(\boldsymbol{\mathcal{X}}) - f_*(\boldsymbol{\mathcal{X}})\,. \tag{3}$$

Then, Lemma 3 shows that the Bernstein condition implies (C_5) when $\mathcal{Z}(f, f_*)$ can be decomposed to the subexponential random variables, $\mathcal{Z}(f, f_*)/\Delta_{f,f_*}$ and Δ_{f,f_*}.

Lemma 3. *Let* $r_0 > 0$ *and suppose that* $(\mu, \ell, \mathcal{F}(r_0), f_*)$ *satisfies the Bernstein condition* (3) *with some* $C > 0$. *Furthermore, suppose that* $\sup_{f \in \mathcal{F}(r_0)} \|\mathcal{Z}(f, f_*)/\Delta_{f,f_*}\|_{\Psi_p} \le R$ *and* $\sup_{f \in \mathcal{F}(r_0)} \|\Delta_{f,f_*}\|_{\Psi_q} \le B$ *with some* $B, R > 0$ *and* $p, q \ge 1$ *having* $\frac{1}{p} + \frac{1}{q} \le 1$. *Then* $(\mu, \ell, \mathcal{F}(r_0), f_*)$ *satisfies* (C_5) *for any* $r \in (0,1)$ *and* $\theta \ge 4tK_n^{2/\min\{p,q\}} \max\big\{\frac{4R^2C}{(t-1)(1-r)}, BR\big\}$, *where* $t > 1$ *is arbitrary, and* $K_n \doteq 4\ln\big(4\min\big\{\sup_{f \in \mathcal{F}(r_0)} \mathbb{K}_0^{1/4}[\Delta_{f,f_*}], \frac{nBR}{r_0}\big\}\big)$.

Proof. Let $\mathcal{Z}_f \doteq \mathcal{Z}(f, f_*)$, $\mathcal{V}_f \doteq \mathcal{Z}(f, f_*)/\Delta_{f,f_*}$, and fix any $f \in \mathcal{F}(r_0)$. Then by the definition of $\mathcal{F}(r_0)$, the Cauchy-Schwartz inequality, and Lemma 12d with $s = 2$, we get

$$\frac{r_0}{n} \le \mathbb{E}[\mathcal{Z}_f] \le \mathbb{E}\big[|\mathcal{V}_f \Delta_{f,f_*}|\big] \le \mathbb{E}\big[\mathcal{V}_f^2\big]^{\frac{1}{2}} \mathbb{E}\big[\Delta_{f,f_*}^2\big]^{\frac{1}{2}} \le 2R\,\mathbb{E}\big[\Delta_{f,f_*}^2\big]^{\frac{1}{2}},$$

which implies $\mathbb{E}\big[\Delta_{f,f_*}^2\big] \ge \big(r_0/(2nR)\big)^2$. Combining this with Lemma 12d for $s = 4$, we obtain $\mathbb{K}_0[\Delta_{f,f_*}] \le \mathbb{E}\big[\Delta_{f,f_*}^4\big]\big(r_0/(2nR)\big)^{-4} \le (4nBR/r_0)^4$. Then, by using Lemma 15 with $\mathcal{W} \leftarrow \Delta_{f,f_*}$, $\mathcal{Z} \leftarrow \mathcal{V}_f$, and $2\ln(64\,\mathbb{K}_0[\Delta_{f,f_*}]) \le 2K_n$, $z \doteq \min\{p, q\}^{-1}$, we get for all $2 \le k \in \mathbb{N}$ that

$$\mathbb{E}\big[|\mathcal{Z}_f|^k\big] = \mathbb{E}\big[|\mathcal{V}_f \Delta_{f,f_*}|^k\big] \le (k!/2)\big(16K_n^{2z}\,\mathbb{E}\big[\Delta_{f,f_*}^2\big]R^2\big)\big(4K_n^{2z}BR\big)^{k-2}.$$

Hence, the conditions of Bernstein's lemma (Lemma 14) hold for $\mathcal{Z}_f$ and $\theta \ge 4tK_n^{2z}BR$ with any $t > 1$, so by $(1 - 4K_n^{2z}BR/\theta)^{-1} \le t/(t-1)$, we obtain

$$\begin{aligned}
\mathbb{E}\big[e^{(r/\theta)\mathbb{E}[\mathcal{Z}_f] - (1/\theta)\mathcal{Z}_f}\big] &= e^{(r-1)\mathbb{E}[\mathcal{Z}_f]/\theta}\,\mathbb{E}\big[e^{(\mathbb{E}[\mathcal{Z}_f] - \mathcal{Z}_f)/\theta}\big] \\
&\le \exp\Big(\frac{r-1}{\theta}\mathbb{E}[\mathcal{Z}_f] + \frac{16tK_n^{2z}R^2}{(t-1)\theta^2}\mathbb{E}[\Delta_{f,f_*}^2]\Big) \\
&\le \exp\left(\frac{\mathbb{E}[\mathcal{Z}_f]}{\theta}\Big(r - 1 + \frac{16tK_n^{2z}R^2C}{(t-1)\theta}\Big)\right) \le 1,
\end{aligned}$$

where in the last step we applied the Bernstein condition (3) also implying $\mathbb{E}[\mathcal{Z}_f] \ge 0$ for all $f \in \mathcal{F}(r_0)$, and used $\theta \ge \frac{16tK_n^{2z}R^2C}{(t-1)(1-r)}$. ∎

Notice that Lemma 3 "splits" the subexponential property of the random variable $\mathcal{Z}(f, f_*)$ between $\mathcal{Z}(f, f_*)/\Delta_{f,f_*}$ and Δ_{f,f_*}. For the squared loss $\ell = \ell_{\text{sq}}$, using $p = q = 2$ provides the sub-Gaussian setting (Lecué and Mendelson, 2013). Furthermore, when the random variable Δ_{f,f_*} is bounded, that is $|\Delta_{f,f_*}| \le B$ a.s., we have $\|\Delta_{f,f_*}\|_{\Psi_q} \le 2B$ for all $q \ge 1$, hence we can use $p = 1$ and $q = \infty$ to cover the setting of uniformly bounded functions and subexponential noise with the squared loss (van de Geer, 2000, Section 9.2). For bounded problems (e.g., Bartlett and Mendelson, 2006; Koltchinskii, 2011), when $\mathcal{Z}(f, f_*)$ is bounded, we can use $p = q = \infty$, and eliminate the K_n term completely.

Lemma 3 also shows that even in the worst case we can set the leading constant in the bound of Theorem 1 as $\theta = O\big(K_n^2 C\big)$ with $K_n^2 = O\big(\ln^2(n/r_0)\big)$, which scales logarithmically in the sample size n, and depends on the regression parameters only through the Bernstein condition (3) of $(\mu, \ell, \mathcal{F}(r_0), f_*)$. In the following sections we investigate this dependence for a few popular regression settings.

2.2.1 Lipschitz losses

Observe that if $\sup_{f \in \mathcal{F}(r_0)} \|\Delta_{f,f_*}\|_{\Psi_q} \le B$ holds for some $B > 0$, then the Bernstein condition is always satisfied for the function class $\mathcal{F}(r_0)$ with any $r_0 > 0$ by $C = 2nB^2/r_0$. To see this, use Lemma 12d with $s = 2$, and $1 < (n/r_0)L_\mu(f, \mathcal{F}_*)$ for any $f \in \mathcal{F}(r_0)$ due to the definition of $\mathcal{F}(r_0)$, to obtain

$$\mathbb{E}\big[\Delta_{f,f_*}^2\big] \le 2B^2 < \frac{2nB^2}{r_0}L_\mu(f, \mathcal{F}_*) = \frac{2nB^2}{r_0}\mathbb{E}[\mathcal{Z}(f, f_*)].$$

When the loss function ℓ is R-Lipschitz in its second argument with some $R > 0$, that is $|\ell(y, \hat{y}_1) - \ell(y, \hat{y}_2)| \le R\,|\hat{y}_1 - \hat{y}_2|$ for all $y, \hat{y}_1, \hat{y}_2 \in \mathbb{R}$, we clearly have $|\mathcal{Z}(f, f_*)| \le R\,|\Delta_{f,f_*}|$

for any $f \in \mathcal{F}$. Then, the requirements of Lemma 3 hold with R, $C = 2nB^2/r_0$, and any $r_0 > 0$, so we obtain (C_5).

As C and so θ scale with n/r_0, this setting is similar to the bounded case (Section 2.1), so we choose $r_0 > 0$ to balance the appropriate terms of Theorem 1. For this, here we use $r_0 = \Theta\big(BR\sqrt{n\,K_n \mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0))}\big)$ which balances $\frac{\theta}{n}\mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0)) = \Theta\big(\frac{K_n(BR)^2}{r_0}\mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0))\big)$ with r_0/n. This way the bound of Theorem 1 scales with $BR\sqrt{K_n \mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0))/n}$, which again cannot be improved in general.[3]

2.2.2 Strongly-convex losses

Now consider a loss function ℓ, which is *η-strongly convex* in its second argument, that is $\ell(y, \lambda\hat{y}_1 + (1-\lambda)\hat{y}_2) \le \lambda\ell(y, \hat{y}_1) + (1-\lambda)\ell(y, \hat{y}_2) - \frac{\eta\lambda(1-\lambda)}{2}|\hat{y}_1 - \hat{y}_2|^2$ holds for all $y, \hat{y}_1, \hat{y}_2 \in \mathbb{R}$ and $\lambda \in (0,1)$. Then, if $R_\mu(f_*) \le R_\mu\big((f+f_*)/2\big)$ is satisfied for all $f \in \mathcal{F}$, the Bernstein condition (3) holds with $C = 4/\eta$. To see this, proceed similarly to Bartlett et al. (2006, Lemma 7) by using the strong convexity property of ℓ to get for all $f \in \mathcal{F}$ that

$$\begin{aligned}\mathbb{E}\big[\Delta_{f,f_*}^2\big] &\le (4/\eta)\Big(R_\mu(f) + R_\mu(f_*) - 2R_\mu\big((f+f_*)/2\big)\Big) \\ &\le (4/\eta)\big(R_\mu(f) - R_\mu(f_*)\big) = (4/\eta)\mathbb{E}\big[\mathcal{Z}(f, f_*)\big]\,.\end{aligned} \tag{4}$$

Notice that the condition $R_\mu(f_*) \le \inf_{f\in\mathcal{F}} R_\mu\big((f+f_*)/2\big)$ is implied by the definition of f_* if either f_* is a regression function defined by the reference class $\mathcal{F}_* = \{\mathbb{X} \to \mathbb{R}\}$, or when $f_* \in \operatorname{argmin}_{f\in\mathcal{F}} R_\mu(f)$ and $\mathcal{F}$ is midpoint convex, that is $f, f_* \in \mathcal{F}$ implies $(f+f_*)/2 \in \mathcal{F}$.

Then, we need $\|\Delta_{f,f_*}\|_{\Psi_q} \le B$ and $\|\mathcal{Z}(f, f_*)/\Delta_{f,f_*}\|_{\Psi_p} \le R < \infty$ to satisfy the requirements of Lemma 3. Again, we get the latter for any Lipschitz loss as in Section 2.2.1. However, here the constant C of the Bernstein condition (3) does not scale with the sample size n, which provides better rates by Theorem 1.

One such example is logistic regression with $\mathcal{Y} \in (0,1)$ a.s. using the cross-entropy loss $\ell = \ell_{\text{ce}}$ where $\ell_{\text{ce}}(y, \hat{y}) \doteq y\ln(y/\hat{y}) + (1-y)\ln\big((1-y)/(1-\hat{y})\big)$ for $y, \hat{y} \in (0,1)$, and a hypothesis class $\mathcal{F} \subseteq \{\mathbb{X} \to [\lambda, 1-\lambda]\}$ with some $\lambda \in (0, 1/2)$. Because the ℓ_{ce} function is $1/\lambda$-Lipschitz and $(1-\lambda)^{-2}$-strongly convex in its second argument over the domain $(0,1) \times [\lambda, 1-\lambda]$, we get the requirements of Lemma 3 with $p = q = \infty$, $B = 1-\lambda$, $R = 1/\lambda$, and $C = 4(1-\lambda)^2$ by (4).[4] Here notice that C does not scale with the sample size n as for the general Lipschitz case in Section 2.2.1, which allows Theorem 1 to deliver better rates.

Notice that the squared loss $\ell = \ell_{\text{sq}}$ is 2-strongly convex, however, it is not Lipschitz over the real line. Fortunately, this is not needed for the condition $\|\mathcal{Z}(f, f_*)/\Delta_{f,f_*}\|_{\Psi_p} \le R$ which holds for some $R > 0$ when $\|\mathcal{Y} - f_*(\boldsymbol{\mathcal{X}})\|_{\Psi_p}$ is bounded due to the decomposition $\mathcal{Z}(f, f_*)/\Delta_{f,f_*} = \Delta_{f,f_*} + 2(f_*(\boldsymbol{\mathcal{X}}) - \mathcal{Y})$.

In the following section, we combine these observations for the squared loss $\ell = \ell_{\text{sq}}$ with Theorem 1 and Lemma 3 to provide a detailed analysis for linear least squares estimation.

3 Linear least squares regression

Here we provide an analysis for the *linear least squares regression* setting, which uses the squared loss and considers ERM estimators over affine hypothesis classes for regression

[3] For example, consider estimating the mean of a standard Gaussian random variable through constant functions using the absolute value loss, and derive the optimal $O(n^{-1/2})$ rate by Theorem 1.

[4] To get these values, use $\partial_z \ell_{\text{ce}}(y, z) = \frac{z-y}{z(1-z)}$ and $\partial_{zz}\ell_{\text{ce}}(y, z) = \frac{y}{z^2} + \frac{1-y}{(1-z)^2}$.

problems with sub-Gaussian distributions defined as

$$\mathbb{M}_{\mathrm{subg}}^{\rho,\sigma,d} \doteq \left\{\mu \,\middle|\, (\boldsymbol{\mathcal{X}},\mathcal{Y}) \sim \mu,\, \boldsymbol{\mathcal{X}} \in \mathbb{R}^d,\, \mathcal{Y} \in \mathbb{R},\, \|\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}}\|_{\Psi_2} \le \rho,\, \|\mathcal{Y} - \mathbb{E}\mathcal{Y}\|_{\Psi_2} \le \sigma\right\},$$

with some sub-Gaussian parameters $\rho, \sigma > 0$, and feature space $\mathbb{X} \doteq \mathbb{R}^d$ with dimension $d \in \mathbb{N}$. Further, we consider affine reference classes $\mathcal{F}_* \subseteq \mathcal{F}_{\mathrm{aff}} \doteq \{\boldsymbol{x} \mapsto \boldsymbol{a}^\top \boldsymbol{x} + b,\, \boldsymbol{x} \in \mathbb{R}^d\}$, and use *least squares estimators* (LSEs), that is ERM estimators (1) using the squared loss $\ell = \ell_{\mathrm{sq}}$, over some hypothesis class within affine functions $\mathcal{F}_n \subseteq \mathcal{F}_{\mathrm{aff}}$.

First, we derive a general result (Corollary 6) which is specialized later for the slope constrained (Section 3.1) and penalized (Section 3.2) settings. For the general result, we set the reference class to the set of slope-bounded affine functions as $\mathcal{F}_* \doteq \mathcal{F}_{\mathrm{aff}}^L$, where $\mathcal{F}_{\mathrm{aff}}^L \doteq \{\boldsymbol{x} \mapsto \boldsymbol{a}^\top \boldsymbol{x} + b : \|\boldsymbol{a}\| \le L\}$ for some Lipschitz bound $L > 0$.

Here we only consider penalty functions which are independent of the bias term satisfying $\frac{\partial}{\partial b}\beta(\boldsymbol{x} \mapsto \boldsymbol{a}^\top \boldsymbol{x} + b) = 0$. Then, we have

$$\overline{\mathcal{Y}} - \boldsymbol{a}^\top \overline{\boldsymbol{\mathcal{X}}} = \operatorname*{argmin}_{b \in \mathbb{R}} \frac{1}{n} \sum_{i=1}^n |\boldsymbol{a}^\top \boldsymbol{\mathcal{X}}_i + b - \mathcal{Y}_i|^2 + \beta_\lambda(\boldsymbol{x} \mapsto \boldsymbol{a}^\top \boldsymbol{x} + b)\,,$$

hence any estimate $f_n \in (\alpha, \beta_\lambda)$-ERM$(\mathcal{F}_{\mathrm{aff}})$ can be expressed as $f_n(\boldsymbol{x}) \doteq \boldsymbol{a}_n^\top(\boldsymbol{x} - \overline{\boldsymbol{\mathcal{X}}}) + \overline{\mathcal{Y}}$ with some $\boldsymbol{a}_n \in \mathbb{R}^d$. Moreover, because $\mathbb{E}[\mathcal{Y} - \boldsymbol{a}^\top \boldsymbol{\mathcal{X}}] = \operatorname{argmin}_{b \in \mathbb{R}} \mathbb{E}\left[|\boldsymbol{a}^\top \boldsymbol{\mathcal{X}} + b - \mathcal{Y}|^2\right]$, we can also write any reference function f_* as $f_*(\boldsymbol{x}) \doteq \boldsymbol{a}_*^\top(\boldsymbol{x} - \mathbb{E}\boldsymbol{\mathcal{X}}) + \mathbb{E}\mathcal{Y}$ with some $\boldsymbol{a}_* \in \mathbb{R}^d$.

Now introduce the following linear function classes:

$$\mathcal{F}_{\mathrm{aff}}^{L,\mu}(t) \doteq \left\{\boldsymbol{x} \mapsto \boldsymbol{a}^\top(\boldsymbol{x} - \mathbb{E}\boldsymbol{\mathcal{X}}) + b : \|\boldsymbol{a}\| \le L,\, b - \mathbb{E}\mathcal{Y} \in [-t, t]\right\},$$
$$\mathcal{F}_{\mathrm{aff}}^{L,n}(t) \doteq \left\{\boldsymbol{x} \mapsto \boldsymbol{a}^\top(\boldsymbol{x} - \overline{\boldsymbol{\mathcal{X}}}) + b : \|\boldsymbol{a}\| \le L,\, b - \overline{\mathcal{Y}} \in [-t, t]\right\},$$

for any $t \ge 0$. Observe that any reference function satisfies $f_* \in \operatorname{argmin}_{f \in \mathcal{F}_{\mathrm{aff}}^{L,\mu}(t)} R_\mu(f)$ for any $t \ge 0$, and any estimate $f_n \in (\alpha, \beta)$-ERM$(\mathcal{F}_{\mathrm{aff}})$ with Lipschitz bound $\|\boldsymbol{a}_n\| \le L$ satisfies $f_n \in (\alpha, \beta)$-ERM$\big(\mathcal{F}_{\mathrm{aff}}^{L,n}(t)\big)$ for all $t \ge 0$.

Because distribution μ is unknown, estimators cannot be represented by the class $\mathcal{F}_{\mathrm{aff}}^{L,\mu}(t)$, just by its data-dependent approximation $\mathcal{F}_{\mathrm{aff}}^{L,n}(t)$. However, as the quantities $\mathbb{E}\boldsymbol{\mathcal{X}}$ and $\mathbb{E}\mathcal{Y}$ are "well-approximated" by $\overline{\boldsymbol{\mathcal{X}}}$ and $\overline{\mathcal{Y}}$ for sub-Gaussian random variables $\boldsymbol{\mathcal{X}}$ and $\mathcal{Y}$, the function classes $\mathcal{F}_{\mathrm{aff}}^{L,\mu}(t)$ and $\mathcal{F}_{\mathrm{aff}}^{L,n}(t)$ are "close". More precisely, one can show (see the proof of Theorem 5) that $\mathcal{F}_{\mathrm{aff}}^{L,\mu}(0) \subseteq \mathcal{F}_{\mathrm{aff}}^{L,n}(t) \subseteq \mathcal{F}_{\mathrm{aff}}^{L,\mu}(2t)$ holds with probability at least $1 - \gamma$, where $t \doteq \Theta\big(\max\{L\rho, \sigma\}\sqrt{\ln(1/\gamma)}\big)$. Hence, Theorem 1 is applicable for such function sets and provides an excess risk upper bound for (α, β)-ERM$(\mathcal{F}_{\mathrm{aff}}^{L,n})$ estimators.

Next we point out that the Lipschitz bound on the slope $\|\boldsymbol{a}_n\|$ can be often improved by using the ERM property (1) when the smallest eigenvalue of the feature covariance matrix is bounded away from zero. Denote the covariance matrix by $\Sigma \doteq \mathbb{E}\left[(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})^\top\right]$ and let its smallest eigenvalue be $\rho^2\eta_\mu \ge 0$. So we have $\Sigma \succeq \rho^2\eta_\mu I_d$, where I_d denotes the $d \times d$ identity matrix. Then consider Lemma 4, which provides a refinement for the Lipschitz bound when $\eta_\mu > 0$.

Lemma 4. *Let $\mu \in \mathbb{M}_{subg}^{\rho,\sigma,d}$ be any distribution such that $\eta_\mu > 0$, $\ell = \ell_{sq}$ be the squared loss, and consider an estimate $f_n \in (\alpha, \beta)$-ERM$(\mathcal{F}_{aff}^{L,n}(0))$ such that $\alpha \le \frac{1}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2$ and $\beta : \mathcal{F}_{aff} \to [0, \infty)$ satisfies $\frac{\partial}{\partial b}\beta(\boldsymbol{x} \mapsto \boldsymbol{a}^\top \boldsymbol{x} + b) = 0$ for some $\gamma \in (0, 1)$. Then, $\mathbb{P}\{\|\boldsymbol{a}_n\| \le \tau(L, \gamma)\} \ge 1 - \gamma$ holds for any $n \in \mathbb{N}$, where*

$$\tau(L, \gamma) \doteq \min\left\{L^2,\, \frac{1}{\eta_\mu}\left(\frac{\sigma^2}{\rho^2} + \frac{dL^2}{n}\right)\tau_{\ln}(d, n, \gamma, \eta_\mu)\right\}^{1/2},$$

and $\tau_{\ln}(d,n,\gamma,\eta_\mu) \doteq 10\big(11\ln(23/\eta_\mu)\ln(3n/\min\{d,n\})+6\big)\ln(6/\gamma)$.

Furthermore, for any $f_* \in \operatorname{argmin}_{f\in\mathcal{F}^L_{aff}} R_\mu(f)$, *we have* $\|\boldsymbol{a}_*\| \le \tau(L,\gamma)$.

Proof. Here we only prove the claim for $\boldsymbol{a}_*$ and provide the proof for $\boldsymbol{a}_n$ later in Appendix B.

Using the definition of f_* and the fact that the constant function $\boldsymbol{x}\mapsto\mathbb{E}\mathcal{Y}$ is in $\mathcal{F}^L_{\text{aff}}$, we have $R_\mu(f_*) \le R_\mu(\boldsymbol{x}\mapsto\mathbb{E}\mathcal{Y})$ which can be rearranged into $\mathbb{E}\big[|\boldsymbol{a}_*^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}})|^2\big] \le 2\mathbb{E}[\boldsymbol{a}_*^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}})(\mathcal{Y}-\mathbb{E}\mathcal{Y})]$. Using this, $2ab\le\frac{a^2}{2}+2b^2$, and Lemma 12d with $s=q=2$, we obtain

$$\rho^2\eta_\mu\|\boldsymbol{a}_*\|^2 \le \boldsymbol{a}_*^\top\Sigma\boldsymbol{a}_* = \mathbb{E}\big[|\boldsymbol{a}_*^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}})|^2\big] \le 4\,\mathbb{E}\big[|\mathcal{Y}-\mathbb{E}\mathcal{Y}|^2\big] \le 4\sigma^2\,,$$

which proves the claim for $\boldsymbol{a}_*$ after rearrangement. ∎

For convenience, define $\tau(L,\gamma)\doteq L$ for distributions μ with $\eta_\mu=0$. Then notice that Lemma 4 improves the Lipschitz bound significantly when η_μ is bounded away from zero, $\sigma/\rho \ll L$, and $n\gg d$.

Finally, we put all the details together and provide the following bound on the excess risk for linear least squares estimation:

Theorem 5. *Consider any distribution* $\mu\in\mathbb{M}^{\rho,\sigma,d}_{subg}$, *the squared loss* $\ell=\ell_{sq}$, *and an estimator* $f_n\in(\alpha,\beta)$*-ERM*$(\mathcal{F}_{aff})$ *with penalty* $\beta:\mathcal{F}_{aff}\to[0,\infty)$ *satisfying* $\frac{\partial}{\partial b}\beta(\boldsymbol{x}\mapsto\boldsymbol{a}^\top\boldsymbol{x}+b)=0$, $\mathbb{P}\{\|\boldsymbol{a}_n\|>L\}\le\gamma/2$ *for some* $L>0$, $\mathbb{P}\big\{\beta(f_*)+\alpha>B_*\big\}\le\gamma/16$ *for some* $\gamma\in(0,1)$ *and reference function* f_*. *Then for all* $n\in\mathbb{N}$, *we have with probability at least* $1-\gamma$ *that*

$$L_\mu\big(f_n,\mathcal{F}^L_{aff}\big) = O\bigg(\frac{d\max\{\tau(L,\gamma)\rho,\sigma\}^2}{n}\,T_{\ln}(n,d,\gamma)+B_*\bigg),$$

where $T_{\ln}(n,d,\gamma)\doteq\big(\ln(e\,n/\min\{d,n\})+\ln\ln(e/\gamma)\big)\ln(e\,n/\min\{d,n\})\ln(1/\gamma)$.

Proof. First, condition on the event $\|\boldsymbol{a}_n\|\le L$, and use *Lemma* 4 with $\gamma\leftarrow\gamma/4$ to get

$$\begin{aligned}\mathbb{P}\{L_\mu(f_n,\mathcal{F}^L_{\text{aff}})>b\} &\le \gamma/2+\mathbb{P}\big\{L_\mu(f_n,\mathcal{F}^L_{\text{aff}})\,\mathbb{I}\{\|\boldsymbol{a}_n\|\le L\}>b\big\}\\ &\le \frac{3\gamma}{4}+\mathbb{P}\big\{L_\mu\big(f_n,\mathcal{F}^L_{\text{aff}}\big)\,\mathbb{I}\{\|\boldsymbol{a}_n\|\le\hat{L}\}>b\big\}\,,\end{aligned}$$

for any $b>0$, where $\hat{L}\doteq\tau(L,\gamma/4)$, and $\mathbb{I}\{\cdot\}$ denotes the indicator function. To find an appropriate b for the second term, we will apply Theorem 1.

Fix $t\doteq t_0\sqrt{\ln(32/\gamma)}$ with $t_0\doteq 2\max\{\hat{L}\rho,\sigma\}$, and define the function sets $\hat{\mathcal{F}}_n\doteq\mathcal{F}^{\hat{L},\mu}_{\text{aff}}(0)$, $\mathcal{F}_n\doteq\mathcal{F}^{\hat{L},n}_{\text{aff}}(t)$, and $\mathcal{F}\doteq\mathcal{F}^{\hat{L},\mu}_{\text{aff}}(2t)$. Notice that $f_n\in(\alpha,\beta)$-ERM$(\mathcal{F}_n)$ holds conditioned on the event $\|\boldsymbol{a}_n\|\le\hat{L}$. Moreover, $f_*\in\hat{\mathcal{F}}_n$ due to Lemma 4, so we have (C_2).

To prove (C_1) for $\hat{\mathcal{F}}_n$, $\mathcal{F}_n$, and $\mathcal{F}$, use $\|\hat{L}\|\overline{\boldsymbol{\mathcal{X}}}-\mathbb{E}\boldsymbol{\mathcal{X}}\|+|\overline{\mathcal{Y}}-\mathbb{E}\mathcal{Y}|\|_{\Psi_2}\le t_0$ due to Lemma 12b, and Lemma 12c, to get

$$\mathbb{P}\big\{\hat{\mathcal{F}}_n\subseteq\mathcal{F}_n\subseteq\mathcal{F}\big\}\ge 1-\mathbb{P}\big\{\hat{L}\|\overline{\boldsymbol{\mathcal{X}}}-\mathbb{E}\boldsymbol{\mathcal{X}}\|+|\overline{\mathcal{Y}}-\mathbb{E}\mathcal{Y}|>t\big\}\ge 1-2e^{-t^2/t_0^2}=1-\gamma/16\,.$$

Next we use Lemma 3 to show (C_5). For this, write $f\in\mathcal{F}$ as $f(\boldsymbol{x})\doteq\boldsymbol{a}^\top(\boldsymbol{x}-\mathbb{E}\boldsymbol{\mathcal{X}})+b$, and observe that

$$|\Delta_{f,f_*}| = |f(\boldsymbol{\mathcal{X}})-f_*(\boldsymbol{\mathcal{X}})| = \big|(\boldsymbol{a}-\boldsymbol{a}_*)^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}})+b-\mathbb{E}\mathcal{Y}\big| \le 2\hat{L}\,\|\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}\|+2t\,.$$

Now pick $f \in \mathcal{F}$ arbitrarily, and use $(a+b)^2 \le 2a^2 + 2b^2$ with Jensen's inequality, to show that Δ_{f,f_*} is sub-Gaussian, that is

$$\mathbb{E}\big[e^{\Delta_{f,f_*}^2/B^2}\big] \le \mathbb{E}\big[e^{8\hat{L}^2\|\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}\|^2/B^2}\big]\, e^{8t^2/B^2} \le 2\,,$$

with $B \doteq 4\max\{\hat{L}\rho, t\}$, so we have $\|\Delta_{f,f_*}\|_{\Psi_2} \le B$. Additionally, as $f_* \in \operatorname{argmin}_{f\in\mathcal{F}} R_\mu(f)$ due to Lemma 4, and the function set $\mathcal{F}$ is convex, the requirements of Lemma 3 are satisfied by (4) and

$$\|\mathcal{Z}(f,f_*)/\Delta_{f,f_*}\|_{\Psi_2} = \|(\boldsymbol{a}+\boldsymbol{a}_*)^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}) + 2(\mathbb{E}[\mathcal{Y}]-\mathcal{Y})\|_{\Psi_2} \le 2(\hat{L}\rho+\sigma) \doteq R\,,$$

so we obtain (C_5) with $r_0 = \min\{d,n\}BR$, $r = 1/2$, and any $\theta = \Omega\big(K_n \max\{B,\sigma\}^2\big)$, where $K_n = O\big(\ln(e\,n/\min\{d,n\})\big)$.

Next, to prove (C_3), write $g \in \mathcal{F}$ as $g(\boldsymbol{x}) \doteq \hat{\boldsymbol{a}}^\top(\boldsymbol{x}-\mathbb{E}\boldsymbol{\mathcal{X}}) + \hat{b}$ with $\|\hat{a}\| \le \hat{L}$ and $|\hat{b}-\mathbb{E}\mathcal{Y}| \le 2t$. Then, using the Cauchy-Schwartz inequality, we obtain

$$\begin{aligned}
\mathcal{Z}(f,g) &= \big(f(\boldsymbol{\mathcal{X}})-g(\boldsymbol{\mathcal{X}})\big)\big(f(\boldsymbol{\mathcal{X}})+g(\boldsymbol{\mathcal{X}})-2\mathcal{Y}\big)\\
&= \big((\boldsymbol{a}+\hat{\boldsymbol{a}})^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}) + b + \hat{b} - 2\mathcal{Y}\big)\big((\boldsymbol{a}-\hat{\boldsymbol{a}})^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}) + b - \hat{b}\big)\\
&\le \Big(2\hat{L}\,\|\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}\| + 4t + 2|\mathcal{Y}-\mathbb{E}\mathcal{Y}|\Big)\begin{bmatrix} 3\hat{L}\,\|\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}\| \\ 4t \end{bmatrix}^\top \begin{bmatrix} \|\boldsymbol{a}-\hat{\boldsymbol{a}}\|/(3\hat{L}) \\ |b-\hat{b}|/(4t) \end{bmatrix}\\
&\le \underbrace{\Big(\big(2\hat{L}\,\|\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}}\| + 4t + 2|\mathcal{Y}-\mathbb{E}\mathcal{Y}|\big)\ln^{-1/2}(32/\gamma)\Big)^2}_{\doteq G(\boldsymbol{\mathcal{X}},\mathcal{Y})}\,\psi(f,g)\,,
\end{aligned}$$

where $\psi(f,g) \doteq \ln(32/\gamma)\big(\|\boldsymbol{a}-\hat{\boldsymbol{a}}\|^2/(2\hat{L})^2 + |b-\hat{b}|^2/(4t)^2\big)^{1/2}$ is a metric on $\mathcal{F}$. As the radius of $\mathcal{F}$ under ψ is bounded by $\ln(32/\gamma)$, that is $\sup_{f\in\mathcal{F}}\psi(f, \boldsymbol{x}\mapsto\mathbb{E}\mathcal{Y}) \le \ln(32/\gamma)$, we have by Lemma 16 that $\mathcal{H}_\psi(\epsilon,\mathcal{F}) \le (d+1)\big(\ln(3/\epsilon) + \ln\ln(32/\gamma)\big)$ for all $\epsilon \in (0,3]$.

Further, as $\|\sqrt{G(\boldsymbol{\mathcal{X}},\mathcal{Y})}\|_{\Psi_2} \le (2t_0+6t)/\sqrt{\ln(32/\gamma)}$, we also have $\|G(\boldsymbol{\mathcal{X}},\mathcal{Y})\|_{\Psi_1} \le (8t_0)^2$, so Lemma 12b implies $\big\|\frac{1}{n}\sum_{i=1}^n G(\boldsymbol{\mathcal{X}}_i,\mathcal{Y}_i)\big\|_{\Psi_1} \le (8t_0)^2$. Hence, due to Lemma 12c, we obtain $\mathbb{P}\big\{\mathbb{E}[G(\boldsymbol{\mathcal{X}},\mathcal{Y})] + \frac{1}{n}\sum_{i=1}^n G(\boldsymbol{\mathcal{X}}_i,\mathcal{Y}_i) > T\big\} \le \gamma/8$ with $T \doteq \mathbb{E}[G(\boldsymbol{\mathcal{X}},\mathcal{Y})] + 64\,t_0^2\ln(32/\gamma)$ having $T < 128\,t^2$ by Lemma 12d with $s=1$ and the definition of t.

Finally, we can apply Theorem 1 with $\gamma \leftarrow \gamma/4$, $\delta \doteq \epsilon$ ignoring (C_4) with $S=\infty$, and choosing $\epsilon \doteq \min\{d,n\}/n$ satisfying $\epsilon \in [0,3]$, to get with probability at least $1-\gamma/4$ that

$$\begin{aligned}
L_\mu(f_n, \mathcal{F}_{\text{aff}}^L)\,\mathbb{I}\{\|\boldsymbol{a}_n\| \le \hat{L}\} &\le 2\left(\frac{\theta\mathcal{H}_\psi(\epsilon,\mathcal{F})}{n} + 16\epsilon\,T + B_*\right) + \frac{r_0}{n}\\
&= O\left(\frac{d\theta}{n}\Big(\ln(e\,n/\min\{d,n\}) + \ln\ln(e/\gamma)\Big) + \frac{t^2}{n} + B_*\right),
\end{aligned}$$

which proves the claim by $\theta = O(\ln(e\,n/\min\{d,n\})\,t^2)$, $t^2 = O\big(\max\{\hat{L}\rho,\sigma\}^2\ln(1/\gamma)\big)$, and $r_0 = O(t^2)$. ■

3.1 Linear least squares with quadratic slope constraint

Here we specialize Theorem 5 to train an affine LSE with $\|\cdot\|$-bounded slope without using any penalty term ($\beta \doteq 0$), and provide the following result:

Corollary 6. *Consider any distribution $\mu \in \mathbb{M}^{\rho,\sigma,d}_{subg}$, the squared loss $\ell = \ell_{sq}$, and an estimate $f_n \in \alpha\text{-ERM}(\mathcal{F}^{L,n}_{aff})$ with any α having $\alpha = O\big(\frac{1}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2/n\big)$ with probability at least $\gamma/16$ for some $\gamma \in (0,1)$. Then for all $n \in \mathbb{N}$, we have with probability at least $1-\gamma$ that*

$$L_\mu(f_n, \mathcal{F}^L_{aff}) = O\left(\frac{d \max\{\tau(L,\gamma)\rho, \sigma\}^2}{n} T_{\ln}(n,d,\gamma)\right).$$

Proof. Notice that

$$\frac{1}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2 = \frac{1}{n}\sum_{i=1}^n |\mathcal{Y}_i - \mathbb{E}\mathcal{Y}|^2 - |\overline{\mathcal{Y}} - \mathbb{E}\mathcal{Y}|^2 \quad \Rightarrow \quad \left\| \frac{1}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2 \right\|_{\Psi_1} \le \sigma^2, \tag{5}$$

so we have $\mathbb{P}\{\alpha > B_*\} \le \gamma/16$ by Lemma 12c for $B_* = O(\frac{\sigma^2 \ln(1/\gamma)}{n})$. Then, the claim follows from Theorem 5 with $\beta = 0$ using $\mathbb{P}\{\beta(f_*) + \alpha > B_*\} = \mathbb{P}\{\alpha > B_*\} \le \gamma/16$. ∎

Notice that Corollary 6 provides an $O(1/n)$ bound on the excess risk up to logarithmic factors for any problem in $\mathbb{M}^{\rho,\sigma,d}_{\text{subg}}$ without any further statistical assumptions such as the bounded magnitude of the kurtosis $\sup_{f \in \mathcal{F}^L_{\text{aff}}} \mathbb{K}_0[\Delta_{f,f_*}]$ as discussed in Section 1.2.

Furthermore, the bound of Corollary 6 after integrated by (2) is comparable to the conjecture of Shamir (2015) stating that ERM estimates achieve optimal expected excess risk up to logarithmic factors for bounded distributions. Our bound is only slightly weaker in general than the conjecture by scaling with $d\max\{\tau(L,\gamma)\rho,\sigma\}^2 \le d\max\{L\rho,\sigma\}^2$ instead of $\max\{(L\rho)^2, d\sigma^2\}$. However, if either $\sigma = \Omega(L\rho)$, or $\eta_\mu = \Omega(1)$ and $n = \Omega(d^2)$, then Corollary 6 matches the bound of the conjecture up to logarithmic factors.

3.2 Linear least squares with quadratic slope penalty

Now we drop the fixed Lipschitz bound on the estimators, and use the reference class $\mathcal{F}_* = \mathcal{F}_{\text{aff}}$. Then, for $f_* \in \operatorname{argmin}_{\mathcal{F}_{\text{aff}}} R_\mu(f)$, we write $f_*(\boldsymbol{x}) = \boldsymbol{a}_*^\top(\boldsymbol{x} - \mathbb{E}\boldsymbol{\mathcal{X}}) + \mathbb{E}\mathcal{Y}$ and set $L_* \doteq \|\boldsymbol{a}_*\|$.[5] Then consider the *ridge regression* (Hoerl and Kennard, 1970) estimate $f_n^\lambda \in (0, \beta_\lambda)\text{-ERM}(\mathcal{F}_{\text{aff}})$ using the quadratic penalty term $\beta_\lambda(\boldsymbol{x} \mapsto \boldsymbol{a}^\top\boldsymbol{x} + b) \doteq \lambda\|\boldsymbol{a}\|^2$ with some $\lambda \ge 0$. This estimator can be also computed in closed-form as $f_n^\lambda(\boldsymbol{x}) \doteq \boldsymbol{a}_{n,\lambda}^\top(\boldsymbol{x} - \overline{\boldsymbol{\mathcal{X}}}) + \overline{\mathcal{Y}}$,

$$\boldsymbol{a}_{n,\lambda} \doteq \left(\lambda I_d + \frac{1}{n}\sum_{i=1}^n (\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})(\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})^\top\right)^{-1}\left(\frac{1}{n}\sum_{i=1}^n (\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})\mathcal{Y}_i\right). \tag{6}$$

It is known that minimizing the empirical risk without any slope restriction (i.e. $\lambda = 0$) might result in infinite expected excess risk (Huang and Szepesvári, 2014, Example 3.5). Moreover, Corollary 6 with $p = 2$ implies (by Lagrangian relaxation and $L \leftarrow L_*$) that for each distribution $\mu \in \mathbb{M}^{\rho,\sigma,d}_{\text{subg}}$, there exists $\lambda \ge 0$ such that the excess risk rate of ridge regression is bounded by $O(1/n)$. In this section, we are interested in choosing λ independently of the parameters of μ.

For this, Lemma 17 with (6) provides an upper bound for the Lipschitz factor of f_n^λ as $\|\boldsymbol{a}_{n,\lambda}\| \le \sqrt{\frac{1}{4n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2/\lambda}$ implying $\mathbb{P}\{\|\boldsymbol{a}_n\|^2 > \sigma^2 \ln(4/\gamma)/(4\lambda)\} \le \gamma/2$ due to (5) and Lemma 12c. Then $L = \max\{L_*^2, \sigma^2\ln(4/\gamma)/(4\lambda)\}^{1/2}$ provides a common Lispchitz bound for f_n^λ, f_*, and $f_* \in \operatorname{argmin}_{f \in \mathcal{F}^L_{\text{aff}}} R_\mu(f)$. Hence, we can apply Theorem 5 and get the following result:

[5] If there are multiple f_*, we choose the one with the smallest slope $\|\boldsymbol{a}_*\|$.

Corollary 7. *Consider any problem $\mu \in \mathbb{M}_{subg}^{\rho,\sigma,d}$, the squared loss $\ell = \ell_{sq}$, and an estimate $f_n^\lambda \in (0, \beta_\lambda)$-ERM($\mathcal{F}_{aff}$) with some $\lambda > 0$. Then for all $\gamma \in (0,1)$, any $n \in \mathbb{N}$, and $L_\lambda \doteq \max\{L_*^2, \sigma^2 \ln(4/\gamma)/(4\lambda)\}^{1/2}$, we have with probability at least $1-\gamma$ that*

$$L_\mu(f_n^\lambda, \mathcal{F}_{aff}) = O\left(\left(\frac{d}{n} \max\left\{\tau(L_\lambda, \gamma)\rho, \sigma\right\}^2 + \lambda L_*^2\right) T_{\ln}(n, d, \gamma)\right).$$

Proof. As explained above, we have $\mathbb{P}\{\|\boldsymbol{a}_n\| > L_\lambda\} \le \gamma/2$. Then, the claim follows directly from Theorem 5 using $L = L_\lambda$, $f_* \in \operatorname{argmin}_{f \in \mathcal{F}_{\text{aff}}^L} R_\mu(f)$, $\frac{\partial}{\partial b}\beta_\lambda(\boldsymbol{x} \mapsto \boldsymbol{a}^\top \boldsymbol{x} + b) = 0$, $\alpha = 0$, and $\beta_\lambda(f_*) = \lambda L_*^2$. ∎

Notice that $\tau(L_\lambda, \gamma) \le \max\left\{L_*^2, \sigma^2 \ln(4/\gamma)/(4\lambda)\right\}^{1/2}$, hence for $n \ge d$ and $\lambda = \sqrt{d/n}$, Corollary 7 upper bounds the excess risk $L_\mu(f_n^\lambda, \mathcal{F}_{\text{aff}})$ with probability at least $1-\gamma$ by $O\big(\max\{L_*, \sigma\}^2 (1+\rho)^2 \sqrt{d/n}\, \ln^3(1/\gamma)\big)$ up to logarithmic factors.

If $\eta_\mu = \Omega(1)$, we can also use $\tau^2(L_\lambda, \gamma) \le \frac{1}{\eta_\mu}\big(\frac{\sigma^2}{\rho^2} + \frac{d \max\{L_*^2, \sigma^2 \ln(4/\gamma)/(4\lambda)\}}{n}\big)\tau_{\ln}(d, n, \gamma, \eta_\mu)$. Then, for $n \ge d$ and $\lambda = d/n$, Corollary 7 upper bounds the excess risk $L_\mu(f_n^\lambda, \mathcal{F}_{\text{aff}})$ with probability at least $1-\gamma$ by $O\big(\max\{L_*, \sigma\}^2 (1+\rho)^2 d/n\, \ln^3(1/\gamma)\big)$ up to logarithmic factors.

Finally, we point out that Corollary 7 is comparable to the result of Hsu et al. (2014, Remarks 4 and 12), which uses similar conditions to prove the same rates in terms of d and n with slightly better constants, but only for the bounded setting when $\|\boldsymbol{\mathcal{X}}\| \le \rho$ a.s. holds.

4 Proof of the upper bound

In this section, we finally prove our main result, Theorem 1, our upper bound on the excess risk of ERM estimators.

The strategy is to "remove" the data-dependence of the hypothesis class $\mathcal{F}_n$ by (C_1), and decompose the excess risk to "supremal" and "approximation" error terms. Then we reduce the former to a general concentration inequality (Theorem 11) with (C_3), (C_4), (C_5), and upper bound the latter using the ERM property (1) and (C_2).

To work with probabilistic arguments, we use the following rule without further notice: $\mathbb{P}\{\mathcal{W} + \mathcal{Z} > t + s\} \le \mathbb{P}\{\mathcal{W} > t\} + \mathbb{P}\{\mathcal{Z} > s\}$ for any random variables $\mathcal{W}, \mathcal{Z}$, and $t, s \in \mathbb{R}$, which holds as a simple corollary of the law of total probability.

Use (C_1), and the definition of $\mathcal{F}(r_0)$ implying $L_\mu(f, \mathcal{F}_*) \le r_0/n$ for all $f \in \mathcal{F} \setminus \mathcal{F}(r_0)$, to get with probability at least $1 - \gamma/4$ that

$$L_\mu(f_n, \mathcal{F}_*) = L_\mu(f_n, \mathcal{F}_*)\, \mathbb{I}\{\hat{\mathcal{F}}_n \subseteq \mathcal{F}_n \subseteq \mathcal{F}\} \le L_\mu(f_n, \mathcal{F}_*)\, \mathbb{I}\big\{\hat{\mathcal{F}}_n \subseteq \mathcal{F}_n,\ f_n \in \mathcal{F}(r_0)\big\} + \frac{r_0}{n}\,.$$

Then notice that the ERM property (1) implies $L_n(f_n, f_*) \le \inf_{f \in \mathcal{F}_n} L_n(f, f_*) + \beta(f) + \alpha$, so we can transform the previous excess risk inequality with $L_\mu(f_n, \mathcal{F}_*) = L_\mu(f_n, f_*)$, and ($C_2$), to get with probability at least $1 - \gamma/2$ that

$$\begin{aligned}
L_\mu(f_n, \mathcal{F}_*) &\le \left(L_\mu(f_n, f_*) - \frac{L_n(f_n, f_*)}{r} + \frac{L_n(f_n, f_*)}{r}\right) \mathbb{I}\{\hat{\mathcal{F}}_n \subseteq \mathcal{F}_n,\ f_n \in \mathcal{F}(r_0)\} + \frac{r_0}{n} \\
&\le \frac{1}{r}\left(\sup_{f \in \mathcal{F}(r_0)} \left\{r L_\mu(f, f_*) - L_n(f, f_*)\right\} + \inf_{f \in \hat{\mathcal{F}}_n} L_n(f, f_*) + \beta(f) + \alpha\right) + \frac{r_0}{n} \qquad (7) \\
&\le \frac{1}{r} \sup_{f \in \mathcal{F}(r_0)} \Gamma(f, \mathcal{D}_n) + \frac{B_*}{r} + \frac{r_0}{n}\,,
\end{aligned}$$

where $\Gamma(f, \mathcal{D}_n) \doteq r L_\mu(f, f_*) - L_n(f, f_*)$. It remains to bound the supremal term, for which we use the general concentration inequality Theorem 11 by showing that its conditions (a), (b), (c) are satisfied by (C_3), (C_4), (C_5), respectively.

Recall that $L_\mu(f, f_*) = \mathbb{E}[\mathcal{Z}(f, f_*)]$ and $L_n(f, f_*) = \frac{1}{n}\sum_{i=1}^n \mathcal{Z}_i(f, f_*)$. Notice that $\Gamma(f, \mathcal{D}_n)$ can be rewritten for any f as $\Gamma(f, \mathcal{D}_n) = r\mathbb{E}\big[\mathcal{Z}(f, f_*)\big] - \frac{1}{n}\sum_{i=1}^n \mathcal{Z}_i(f, f_*)$. Define $\Lambda(f, \mathcal{D}_n) \doteq \Gamma(f, \mathcal{D}_n) - \mathbb{E}[\Gamma(f, \mathcal{D}_n)]$, and observe that $\mathbb{E}[\Gamma(f, \mathcal{D}_n)] = (r-1)\mathbb{E}[\mathcal{Z}(f, f_*)]$, so $\Lambda(f, \mathcal{D}_n) = \mathbb{E}\big[\mathcal{Z}(f, f_*)\big] - \frac{1}{n}\sum_{i=1}^n \mathcal{Z}_i(f, f_*)$. Then, we also have for all $f, g \in \mathcal{F}(r_0)$ that

$$\Lambda(f, \mathcal{D}_n) - \Lambda(g, \mathcal{D}_n) = \mathbb{E}\big[\mathcal{Z}(f, g)\big] - \frac{1}{n}\sum_{i=1}^n \mathcal{Z}_i(f, g)\,. \tag{8}$$

Because $\Lambda(f, \mathcal{D}_n) - \Lambda(g, \mathcal{D}_n)$ is the sum of n independent, centered random variables $\mathbb{E}[\mathcal{Z}(f, g)] - \mathcal{Z}_i(f, g)$ with $\|\mathbb{E}[\mathcal{Z}(f, g)] - \mathcal{Z}_i(f, g)\|_{\Psi_q} \le S\,\psi(f, g)$ according to (C_4), Lemma 13 implies that $\|\Lambda(f, \mathcal{D}_n) - \Lambda(g, \mathcal{D}_n)\|_{\Psi_q} \le (4S\,n^{-1/2})\,\psi(f, g)$ for any $f, g \in \mathcal{F}(r_0)$. Hence, Λ and ψ satisfy Theorem 11b with $\mathcal{W} \leftarrow \mathcal{D}_n$, $S \leftarrow 4S\,n^{-1/2}$, and $\mathcal{F} \leftarrow \mathcal{F}(r_0)$.

Next, using (C_3), we can upper bound (8) a.s. by $\tau(\mathcal{D}_n)\,\psi(f, g)$, where

$$\tau(\mathcal{D}_n) \doteq \mathbb{E}[G(\boldsymbol{\mathcal{X}}, \mathcal{Y})] + \frac{1}{n}\sum_{i=1}^n G(\boldsymbol{\mathcal{X}}_i, \mathcal{Y}_i)\,.$$

Hence, Λ, ψ, and τ satisfies Theorem 11a with T as given by (C_3) and $\gamma \leftarrow \gamma/2$.

Finally, using the i.i.d. property of the sample $\mathcal{D}_n$ and (C_5), we have for any $f \in \mathcal{F}(r_0)$ that

$$\mathbb{E}\Big[e^{\Gamma(f, \mathcal{D}_n)/(\theta/n)}\Big] = \prod_{i=1}^n \mathbb{E}\left[e^{(r/\theta)\mathbb{E}[\mathcal{Z}(f, f_*)] - (1/\theta)\mathcal{Z}_i(f, f_*)}\right] \le 1\,. \tag{9}$$

So Γ satisfies Theorem 11c with $\theta \leftarrow \theta/n$.

Hence, all the requirements of Theorem 11 hold, and we get with probability at least $1 - \gamma/2$ that

$$\sup_{f \in \mathcal{F}(r_0)} \Gamma(f, \mathcal{D}_n) \le \theta\frac{\mathcal{H}_\psi(\epsilon, \mathcal{F}(r_0)) + \ln(\frac{4}{\gamma})}{n} + \frac{32S}{\sqrt{n}}\int_\delta^\epsilon \sqrt[q]{2\mathcal{H}_\psi(z, \mathcal{F}(r_0)) + \ln\big(\frac{32\,\epsilon}{z\gamma}\big)}\,dz + 8\delta T\,.$$

Combining this with (7) proves the claim.

4.1 Suprema of empirical processes

In this section, our goal is to prove Theorem 11, which we used in Section 4 as the main tool to prove Theorem 1. For this, we start with finite class lemmas, then adapt the classical chaining argument (e.g., Pollard, 1990, Section 3; van de Geer, 2000, Chapter 3; Boucheron et al., 2012, Section 13.1) to our setting, and finally put these together to prove Theorem 11.

First, consider the probabilistic version of the well-known inequality about the maximum of finitely many random variables (e.g., Cesa-Bianchi and Lugosi, 1999, Lemma 7; Boucheron et al., 2012, Theorem 2.5).

Lemma 8. *Let $\mathcal{F}$ be a nonempty, finite set (that is $1 \le |\mathcal{F}| < \infty$), $\sigma \in [0, \infty)$, and $\mathcal{W}_f$ be random variables such that $\|\mathcal{W}_f\|_{\Psi_q} \le \sigma$ holds for all $f \in \mathcal{F}$. Then for all $\gamma > 0$, $\mathbb{P}\big\{\max_{f \in \mathcal{F}} \mathcal{W}_f > \sigma\sqrt[q]{\ln(2|\mathcal{F}|/\gamma)}\big\} \le \gamma$.*

Proof. The claim is trivial for $\sigma = 0$. Let $\sigma > 0$ and set $t \doteq \sigma \sqrt[q]{\ln(2|\mathcal{F}|/\gamma)}$. Then, using the union bound and Lemma 12c, we get

$$\mathbb{P}\Big\{ \max_{f\in\mathcal{F}} \mathcal{W}_f > t \Big\} \le \sum_{f\in\mathcal{F}} \mathbb{P}\big\{\mathcal{W}_f > t\big\} \le 2|\mathcal{F}| e^{-t^q/\sigma^q} = \gamma\,. \qquad \blacksquare$$

When a moment condition, similar to (C_5), is satisfied for $\mathcal{W}_f$, Lemma 8 can be strengthened by the following result (for further explanation, see the discussion after Lemma 9).

Lemma 9. *Let $\mathcal{F}$ be a nonempty, finite set (that is $1 \le |\mathcal{F}| < \infty$), $\theta \in (0,\infty)$, and $\mathcal{W}_f$ be random variables such that $\mathbb{E}\big[e^{\mathcal{W}_f/\theta}\big] \le 1$ holds for all $f \in \mathcal{F}$. Then for all $\gamma > 0$, $\mathbb{P}\big\{ \max_{f\in\mathcal{F}} \mathcal{W}_f > \theta \ln(|\mathcal{F}|/\gamma)\big\} \le \gamma$.*

Proof. Set $t \doteq \theta \ln(|\mathcal{F}|/\gamma)$. Then, using the union and Chernoff bounds, we get

$$\mathbb{P}\Big\{ \max_{f\in\mathcal{F}} \mathcal{W}_f > t \Big\} \le \sum_{f\in\mathcal{F}} \mathbb{P}\big\{\mathcal{W}_f > t\big\} \le |\mathcal{F}|\, e^{-t/\theta}\, \mathbb{E}\big[e^{\mathcal{W}_f/\theta}\big] \le |\mathcal{F}|\, e^{-t/\theta} = \gamma\,. \qquad \blacksquare$$

To see that Lemma 9 is indeed stronger than Lemma 8 for our purposes, notice that they scale differently in their parameters θ and σ when applied to averages of independent random variables. If $\mathcal{W}_f^{(1)}, \dots, \mathcal{W}_f^{(n)}$ are $n \in \mathbb{N}$ independent centered random variables with $\big\|\mathcal{W}_f^{(i)}\big\|_{\Psi_q} \le \sigma$ for $q \in \{1,2\}$, their average satisfies $\Big\|\frac{1}{n}\sum_{i=1}^n \mathcal{W}_p^{(i)}\Big\|_{\Psi_q} \le 4\sigma/\sqrt{n}$ by Lemma 13. On the other hand, it is straightforward to show (as we did by Equation 9) that if $\mathcal{W}_f^{(1)}, \dots, \mathcal{W}_f^{(n)}$ are n independent (not necessarily centered) random variables with $\mathbb{E}\big[\exp(\mathcal{W}_f^{(i)}/\theta)\big] \le 1$, then their average satisfies the moment condition with θ/n. This speed-up, from $\frac{\sigma}{\sqrt{n}} \ln|\mathcal{F}|$ to $\frac{\theta}{n}\ln|\mathcal{F}|$, will allow us to derive better bounds when the moment condition (C_5) holds with θ being independent of n.

We now extend Lemma 8 to infinite classes by a probabilistic version of the standard chaining argument. The proof goes along the development of Lemma 3.4 of Pollard (1990), replacing the packing sets by internal covering numbers (for better numerical constants) and the sample continuity condition by uniform Lipschitzness (for truncating the integral at δ). The result is also similar to Proposition 3 of Cesa-Bianchi and Lugosi (1999), which works for the sub-Gaussian case, provides and expected value result, and uses external covering numbers with a slightly different chaining argument.

Theorem 10. *Let $(\mathcal{F},\psi)$ be a separable metric space, $\mathcal{W}$ be a random variable on some set $\mathbb{W}$, and $\Lambda : \mathcal{F}\times\mathcal{F} \to \mathbb{R}$ be a function. Furthermore, suppose that the following conditions hold:*

- *(a) there exists $\gamma \in (0,1)$, $T \ge 0$, and $\tau : \mathbb{W} \to [0,\infty)$ such that $\mathbb{P}\{\tau(\mathcal{W}) > T\} \le \gamma/2$, and $\Lambda(f,\mathcal{W}) - \Lambda(g,\mathcal{W}) \le \psi(f,g)\,\tau(\mathcal{W})$ a.s. for all $f,g \in \mathcal{F}$,*
- *(b) $\|\Lambda(f,\mathcal{W}) - \Lambda(g,\mathcal{W})\|_{\Psi_q} \le S\,\psi(f,g)$ with some $S \in [0,\infty]$ for all $f,g \in \mathcal{F}$,*
- *(c) there exist $\beta \ge 0$ and $f_0 \in \mathcal{F}$ such that $\Lambda(f_0,\mathcal{W}) = 0$ a.s., and $\sup_{f\in\mathcal{F}} \psi(f,f_0) \le \beta$.*

Then, for all $0 \le \delta \le \beta/2$, we have with probability at least $1-\gamma$ that

$$\sup_{f\in\mathcal{F}} \Lambda(f,W) \le 4S \int_\delta^{\beta/2} \sqrt[q]{\mathcal{H}_\psi(z,\mathcal{F}) + \ln(4\beta/(z\gamma))}\, dz + 4\delta T\,.$$

Furthermore, the result holds without (b), that is using $S = \infty$, $\delta = \beta/2$, and $\infty\cdot 0 = 0$.

Proof. If there exists $z \in (\delta, \beta/2]$ such that $\mathcal{N}_\psi(z, \mathcal{F}) = \infty$, then the integral is infinite and so the claim is trivial. The claim is also trivial for $\beta = 0$ or $S = 0$. Now assume that $0 < \beta$, $S \in (0, \infty)$ and $\mathcal{N}_\psi(z, \mathcal{F}) < \infty$ for all $z \in (\delta, \beta/2]$.

Let $\delta > 0$ and $m \in \mathbb{N} \cup \{0\}$ be such that $2\delta \leq \beta/2^m < 4\delta$. Further, let $\mathcal{F}_0 \doteq \{f_0\}$, $\epsilon_0 \doteq \beta$, $\epsilon_k \doteq \beta/2^k$, and $\mathcal{F}_k$ be an ϵ_k-cover of $\mathcal{F}$ under ψ with minimal cardinality for all $k \in \{1, \dots, m\}$. Notice that $\mathcal{F}_0$ is an ϵ_0-cover by (c), and define $g_k(f)$ to be the closest element to $f \in \mathcal{F}$ in $\mathcal{F}_k$ for all $k = 0, \dots, m$, that is $g_k(f) \in \operatorname{argmin}_{g \in \mathcal{F}_k} \psi(f, g)$.

Fix some $k \in \{0, \dots, m-1\}$ and $f \in \mathcal{F}_{k+1}$. When $k = 0$, we have $\psi(f, g_k(f)) = \psi(f_0, f) \leq \beta = \epsilon_0$ by (c), while for $k > 0$, $\psi(f, g_k(f)) \leq \epsilon_k$ holds by the definition of $\mathcal{F}_k$. So (b) implies that $\|\Lambda(f, \mathcal{W}) - \Lambda(g_k(f), \mathcal{W})\|_{\Psi_q} \leq \epsilon_k S$. Then, we can chain maximal inequalities for all $k = 0, \dots, m-1$ by using $\mathbb{P}\{\mathcal{X} + \mathcal{Y} > t\} \leq \mathbb{P}\{\mathcal{X} > t - t_0\} + \mathbb{P}\{\mathcal{Y} > t_0\}$ for any random variables $\mathcal{X}, \mathcal{Y}$, and Lemma 8 with $b_k \doteq \epsilon_k S \sqrt[q]{\ln(2|\mathcal{F}_{k+1}|/\gamma_k)}$ and $\gamma_k \doteq 2^{-(k+2)}\gamma$, to get for all $t_k > 0$ that

$$\begin{aligned}
\mathbb{P}\Big\{ \max_{f \in \mathcal{F}_{k+1}} \Lambda(f, \mathcal{W}) > t_k \Big\} &= \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_{k+1}} \Lambda(g_k(f), \mathcal{W}) + \Lambda(f, \mathcal{W}) - \Lambda(g_k(f), \mathcal{W}) > t_k \Big\} \\
&\leq \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_k} \Lambda(f, \mathcal{W}) > t_k - b_k \Big\} + \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_{k+1}} \Lambda(f, \mathcal{W}) - \Lambda(g_k(f), \mathcal{W}) > b_k \Big\} \\
&\leq \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_k} \Lambda(f, \mathcal{W}) > t_k - b_k \Big\} + \gamma_k \,.
\end{aligned} \tag{10}$$

Additionally, using (a), $\sup_{f \in \mathcal{F}} \psi(f, g_m(f)) \leq \epsilon_m < 4\delta$, and $\mathbb{P}\{\tau(\mathcal{W}) > T\} \leq \gamma/2$, we get for all $t > 0$ that

$$\begin{aligned}
\mathbb{P}\Big\{ \sup_{f \in \mathcal{F}} \Lambda(f, \mathcal{W}) > t \Big\} &= \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_m} \Lambda(g_m(f), \mathcal{W}) + \Lambda(f, \mathcal{W}) - \Lambda(g_m(f), \mathcal{W}) > t \Big\} \\
&\leq \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_m} \Lambda(f, \mathcal{W}) > t - 4\delta T \Big\} + \mathbb{P}\Big\{ \sup_{f \in \mathcal{F}} \Lambda(f, \mathcal{W}) - \Lambda(g_m(f), \mathcal{W}) > 4\delta T \Big\} \\
&\leq \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_m} \Lambda(f, \mathcal{W}) > t - 4\delta T \Big\} + \mathbb{P}\Big\{ \sup_{f \in \mathcal{F}} \psi(f, g_m(f))\, \tau(\mathcal{W}) > 4\delta T \Big\} \\
&\leq \mathbb{P}\Big\{ \max_{f \in \mathcal{F}_m} \Lambda(f, \mathcal{W}) > t - 4\delta T \Big\} + \gamma/2 \,.
\end{aligned} \tag{11}$$

Then, set $t \doteq 4\delta T + \sum_{k=0}^{m-1} b_k$, and use (10) repeatedly with $k = m-1, m-2, \dots, 0$ and $t_k \doteq \sum_{l=0}^{k} b_l$ to bound the first term of (11) with $\mathbb{P}\{\max_{f \in \mathcal{F}_0} \Lambda(f, \mathcal{W}) > b_0\} = 0$ and $\sum_{k=0}^{m-1} \gamma_k = \frac{\gamma}{4} \sum_{k=0}^{m-1} 2^{-k} < \gamma/2$, which provides with probability at least $1 - \gamma$ that

$$\begin{aligned}
\sup_{f \in \mathcal{F}} \Lambda(f, \mathcal{W}) &\leq 4\delta T + S \sum_{k=0}^{m-1} \epsilon_k \sqrt[q]{\ln(2|\mathcal{F}_{k+1}|/\gamma_k)} \\
&= 4\delta T + 4S \sum_{k=0}^{m-1} \epsilon_{k+2} \sqrt[q]{\mathcal{H}_\psi(\epsilon_{k+1}, \mathcal{F}) + \ln(4\beta/(\gamma\epsilon_{k+1}))} \,,
\end{aligned} \tag{12}$$

because $\epsilon_k = 4\epsilon_{k+2}$, $\ln|\mathcal{F}_{k+1}| = \mathcal{H}_\psi(\epsilon_{k+1}, \mathcal{F})$, and $2/\gamma_k = 4\beta/(\gamma\epsilon_{k+1})$.

Next, to upper bound the sum of (12) by an integral, use the nondecreasing property of $\mathcal{H}_\psi(z, \mathcal{F})$ as $z \to 0$, and $\epsilon_{k+1} - \epsilon_{k+2} = \epsilon_{k+2}$ to obtain

$$\epsilon_{k+2} \sqrt[q]{\mathcal{H}_\psi(\epsilon_{k+1}, \mathcal{F}) + \ln(4\beta/(\gamma\epsilon_{k+1}))} \leq \int_{\epsilon_{k+2}}^{\epsilon_{k+1}} \sqrt[q]{\mathcal{H}_\psi(z, \mathcal{F}) + \ln(4\beta/(z\gamma))}\, dz \,,$$

for all $k = 0, \dots, m-1$ with $\epsilon_{m+1} \doteq \beta/2^{m+1}$. Plugging this into (12), we get the claim by $\epsilon_1 = \beta/2$ and $\epsilon_{m+1} \geq \delta$.

Finally notice that for $\delta = \beta/2$ (that is $m = 0$), we use only (11) and ignore (b) altogether, hence justifying the $0 \cdot \infty = 0$ convention for the $S = \infty$ case. Furthermore, the $\delta = 0$ case can be obtained through the limit $\delta \to 0$. ■

Now we extend the improved finite class lemma (Lemma 9) to infinite classes and prove the main result of this section, Theorem 11. The idea behind the proof is to apply Lemma 9 in the first step of the chain and the previously developed chaining technique (Theorem 10) to the remainder.

Theorem 11. *Let $(\mathcal{F}, \psi)$ be a separable metric space, $\mathcal{W}$ be a random variable on some set $\mathbb{W}$, and $\Gamma : \mathcal{F} \times \mathcal{F} \to \mathbb{R}$ be a function. Furthermore, define $\Lambda(f, w) \doteq \Gamma(f, w) - \mathbb{E}[\Gamma(f, \mathcal{W})]$ for all $f \in \mathcal{F}$, $w \in \mathbb{W}$, and suppose that the following conditions hold:*

(a) there exists $\gamma \in (0, 1)$, $T \geq 0$, and $\tau : \mathbb{W} \to [0, \infty)$ such that $\mathbb{P}\{\tau(\mathcal{W}) > T\} \leq \gamma/2$, and $\Lambda(f, \mathcal{W}) - \Lambda(g, \mathcal{W}) \leq \psi(f, g)\, \tau(\mathcal{W})$ a.s. for all $f, g \in \mathcal{F}$,

(b) $\|\Lambda(f, \mathcal{W}) - \Lambda(g, \mathcal{W})\|_{\Psi_q} \leq S\, \psi(f, g)$ with some $S \in [0, \infty]$ for all $f, g \in \mathcal{F}$,

(c) there exists $\theta > 0$ such that $\mathbb{E}\big[\exp\big(\Gamma(f, \mathcal{W})/\theta\big)\big] \leq 1$ for all $f \in \mathcal{F}$.

Then, for all $0 \leq \delta \leq \epsilon$, we have with probability at least $1 - \gamma$ that

$$\sup_{f \in \mathcal{F}} \Gamma(f, \mathcal{W}) \leq \theta\big(\mathcal{H}_\psi(\epsilon, \mathcal{F}) + \ln(2/\gamma)\big) + 8S \int_\delta^\epsilon \sqrt[q]{2\, \mathcal{H}_\psi(z, \mathcal{F}) + \ln(16\epsilon/(z\gamma))}\, dz + 8\delta\, T\,.$$

Furthermore, the result holds without (b), that is using $S = \infty$, $\delta = \beta/2$, and $\infty \cdot 0 = 0$.

Proof. Fix $0 < \delta \leq \epsilon$. When $\mathcal{N}_\psi(z, \mathcal{F}) = \infty$ for some $z \in (\delta, \epsilon]$, the claim is trivial, so we can assume that $\mathcal{N}_\psi(z, \mathcal{F}) < \infty$ for all $z \in (\delta, \epsilon]$.

Let $\mathcal{F}_\epsilon$ be an ϵ-cover of $\mathcal{F}$ under ψ with minimal cardinality and define g_f to be the closest element to $f \in \mathcal{F}$ in $\mathcal{F}_\epsilon$, that is $g_f \in \operatorname{argmin}_{g \in \mathcal{F}_\epsilon} \psi(f, g)$. Due to Jensen's inequality and (c), we have $\mathbb{E}[\Gamma(f, \mathcal{W})] \leq 0$ for all $f \in \mathcal{F}$. Define $g_f^* \in \operatorname{argmax}_{g \in \mathcal{F} : \psi(g, g_f) \leq \epsilon} \mathbb{E}[\Gamma(g, \mathcal{W})]$.[6] Then, due to $\psi(g_f, f) \leq \epsilon$, we have $\mathbb{E}[\Gamma(f, \mathcal{W})] \leq \mathbb{E}[\Gamma(g_f^*, \mathcal{W})]$ for all $f \in \mathcal{F}$. Further, $\psi(f, g_f^*) \leq \psi(f, g_f) + \psi(g_f, g_f^*) \leq 2\epsilon$ so $\mathcal{F}_\epsilon^* \doteq \{g_f^* : g_f \in \mathcal{F}_\epsilon\}$ is a 2ϵ-cover of $\mathcal{F}$ under ψ with $|\mathcal{F}_\epsilon^*| = |\mathcal{F}_\epsilon| = \mathcal{N}_\psi(\epsilon, \mathcal{F})$.

Now, for the first step of the chain, consider the following decomposition,

$$\begin{aligned}
\sup_{f \in \mathcal{F}} \Gamma(f, \mathcal{W}) &= \sup_{f \in \mathcal{F}} \Big\{ \Gamma(g_f^*, \mathcal{W}) + \Gamma(f, \mathcal{W}) - \Gamma(g_f^*, \mathcal{W}) \Big\} \\
&= \sup_{f \in \mathcal{F}} \Big\{ \Gamma(g_f^*, \mathcal{W}) + \Lambda(f, \mathcal{W}) - \Lambda(g_f^*, \mathcal{W}) + \mathbb{E}\big[\Gamma(f, \mathcal{W}) - \Gamma(g_f^*, \mathcal{W})\big] \Big\} \qquad (13) \\
&\leq \max_{g \in \mathcal{F}_\epsilon^*} \Gamma(g, \mathcal{W}) + \sup_{f \in \mathcal{F}} \Big\{ \Lambda(f, \mathcal{W}) - \Lambda\big(g_f^*, \mathcal{W}\big) \Big\}\,.
\end{aligned}$$

Then by Lemma 9 and (c), we obtain with probability at least $1 - \gamma/2$ that

$$\max_{g \in \mathcal{F}_\epsilon^*} \Gamma(g, \mathcal{W}) \leq \theta \ln\big(2|\mathcal{F}_\epsilon^*|/\gamma\big) = \theta\big(\mathcal{H}_\psi(\epsilon, \mathcal{F}) + \ln(2/\gamma)\big)\,. \qquad (14)$$

[6] If such g_f^* element does not exist, one can choose another element which is arbitrary close to the supremum and shrink the gap to zero at the end of the analysis.

The rest of the proof is about to upper bound the supremal term on the right side of (13), that is $\sup_{f\in\mathcal{F}} \Lambda(f,\mathcal{W}) - \Lambda(g_f^*,\mathcal{W})$, by using the chaining result of Theorem 10.

Let $\mathcal{K} \doteq \{(f, g_f^*) : f \in \mathcal{F}\} \subseteq \mathcal{F}\times\mathcal{F}_\epsilon^*$ and choose $f_0 \in \operatorname{argmax}_{f\in\mathcal{F}_\epsilon^*} \mathbb{E}[\Gamma(f,\mathcal{W})]$ so that $f_0 = g_{f_0}^*$ (since $\psi(g_{f_0}, f_0) \le \epsilon$), implying $(f_0, f_0) \in \mathcal{K}$. Additionally, define

$$\tilde{\Lambda}\big((f,g_f^*),w\big) \doteq \Lambda(f,w) - \Lambda(g_f^*,w)\,,$$
$$\tilde{\psi}\big((f,g_f^*),(h,g_h^*)\big) \doteq \min\big\{\psi(f,h) + \psi(g_f^*,g_h^*),\, 4\epsilon\big\}\,,$$

for all $(f,g_f^*),(h,g_h^*) \in \mathcal{F}\times\mathcal{F}_\epsilon^*$, and $w \in \mathbb{W}$. Now notice that $(\mathcal{F}\times\mathcal{F}_\epsilon^*, \tilde{\psi})$ is a metric space,[7] and for any $f \in \mathcal{F}$ we have that

$$\tilde{\Lambda}\big((f_0,f_0),\mathcal{W}\big) = 0 \text{ a.s.}\,, \quad \tilde{\psi}\big((f_0,f_0),(f,g_f^*)\big) \le 4\epsilon\,,$$

hence, $\tilde{\psi}$ and $(f_0,f_0) \in \mathcal{K}$ satisfies Theorem 10c with $\beta = 4\epsilon$.

Further, since $\psi(f,g_f^*) \le 2\epsilon$ holds for all $(f,g_f^*) \in \mathcal{K}$, (b) and Lemma 12b with $q \ge 1$ implies for all $(f,g_f^*),(h,g_h^*) \in \mathcal{K}$ that

$$\begin{aligned}
\|\tilde{\Lambda}((f,g_f^*),\mathcal{W}) - \tilde{\Lambda}((h,g_h^*),\mathcal{W})\|_{\Psi_q} &\le \|\Lambda(f,\mathcal{W}) - \Lambda(h,\mathcal{W})\|_{\Psi_q} + \|\Lambda(g_h^*,\mathcal{W}) - \Lambda(g_f^*,\mathcal{W})\|_{\Psi_q} \\
&\le \big(\psi(f,h) + \psi(g_h^*,g_f^*)\big)\, S\,, \\
\|\tilde{\Lambda}((f,g_f^*),\mathcal{W}) - \tilde{\Lambda}((h,g_h^*),\mathcal{W})\|_{\Psi_q} &\le \|\Lambda(f,\mathcal{W}) - \Lambda(g_f^*,\mathcal{W})\|_{\Psi_q} + \|\Lambda(g_h^*,\mathcal{W}) - \Lambda(h,\mathcal{W})\|_{\Psi_q} \\
&\le 4\epsilon\, S\,,
\end{aligned}$$

hence, $\tilde{\Lambda}$ and $\tilde{\psi}$ satisfies Theorem 10b with S.

Similarly, (a) implies that

$$\begin{aligned}
\tilde{\Lambda}\big((f,g_f^*),\mathcal{W}\big) - \tilde{\Lambda}\big((h,g_h^*),\mathcal{W}\big) &= \Lambda(f,\mathcal{W}) - \Lambda(h,\mathcal{W}) + \Lambda(g_h^*,\mathcal{W}) - \Lambda(g_f^*,\mathcal{W}) \\
&\le \big(\psi(f,h) + \psi(g_h^*,g_f^*)\big)\,\tau(\mathcal{W}) \text{ a.s.}\,, \\
\tilde{\Lambda}\big((f,g_f^*),\mathcal{W}\big) - \tilde{\Lambda}\big((h,g_h^*),\mathcal{W}\big) &= \Lambda(f,\mathcal{W}) - \Lambda(g_f^*,\mathcal{W}) + \Lambda(g_h^*,\mathcal{W}) - \Lambda(h,\mathcal{W}) \\
&\le \big(\psi(f,g_f^*) + \psi(g_h^*,h)\big)\,\tau(\mathcal{W}) \le 4\epsilon\,\tau(\mathcal{W}) \text{ a.s.}\,,
\end{aligned}$$

so $\tilde{\Lambda}$ and $\tilde{\psi}$ satisfies Theorem 10a with τ.

Then the requirements of Theorem 10 hold (using $\mathcal{F} \leftarrow \mathcal{K}$, $\Lambda \leftarrow \tilde{\Lambda}$, $\psi \leftarrow \tilde{\psi}$, $f_0 \leftarrow (f_0,f_0)$, $\beta = 4\epsilon$, $\delta \leftarrow 2\delta$, $\gamma \leftarrow \gamma/2$), so we get with probability at least $1-\gamma/2$ that

$$\begin{aligned}
\sup_{f\in\mathcal{F}} \big\{\Lambda(f,\mathcal{W}) - \Lambda\big(g_f^*,\mathcal{W}\big)\big\} &= \sup_{\kappa\in\mathcal{K}} \tilde{\Lambda}(\kappa,\mathcal{W}) \\
&\le 4S \int_{2\delta}^{2\epsilon} \sqrt[q]{\mathcal{H}_{\tilde{\psi}}(z,\mathcal{K}) + \ln(32\epsilon/(z\gamma))}\, dz + 8\delta\, T \qquad (15) \\
&= 8S \int_{\delta}^{\epsilon} \sqrt[q]{\mathcal{H}_{\tilde{\psi}}(2z,\mathcal{K}) + \ln(16\epsilon/(z\gamma))}\, dz + 8\delta\, T\,.
\end{aligned}$$

It remains to bound the entropy of $(\mathcal{K},\tilde{\psi})$. For any $z \in (\delta,\epsilon]$, let $\mathcal{F}_z$ be a z-cover of $\mathcal{F}$ under ψ with minimal cardinality and define $\mathcal{K}_z \doteq \mathcal{F}_z \times \mathcal{F}_\epsilon^*$. Then $\mathcal{K}_z$ is an external z-cover of $\mathcal{K}$ in the metric space $(\mathcal{F}\times\mathcal{F}_\epsilon^*, \tilde{\psi})$, which means that $\mathcal{K}_z$ might not be a subset of $\mathcal{K}$, but for any $\kappa \in \mathcal{K}$ there exists $\hat{\kappa} \in \mathcal{K}_z$ for which $\tilde{\psi}(\kappa,\hat{\kappa}) \le z$. Then, as $|\mathcal{K}_z| = |\mathcal{F}_z| \cdot |\mathcal{F}_\epsilon^*| \le \mathcal{N}_\psi(z,\mathcal{F})^2$, using the relation between internal and external covering numbers (Dudley, 1999, Theorem 1.2.1), we get $\ln \mathcal{N}_{\tilde{\psi}}(2z,\mathcal{K}) \le \ln|\mathcal{K}_z| \le 2\ln \mathcal{N}_\psi(z,\mathcal{F})$. Finally, combining this with (15), (14), and (13), we get the claim.

Finally, obtain the $\delta = 0$ case as well through the limit $\delta \to 0$. ■

[7] To prove the triangle inequality, use $\min\{a+b,c\} \le \min\{a,c\} + \min\{b,c\}$ for $a,b,c \ge 0$.

5 Conclusion

In this paper we provided a probabilistic excess risk upper bound for ERM estimators in the random design setting. Although we demonstrated the strength of the result for LSEs, the result is applicable beyond the squared loss and linear models.

In fact, Theorem 1 can be used to recover many nonparametric results in the literature, for example the convex regression result of Lim (2014, Theorem 1) for max-affine estimators with n hyperplanes, or its sieved variant by using only $\lceil n^{d/(d+4)} \rceil$ hyperplanes as discussed by Balázs et al. (2015, Theorem 4.2). What is interesting though that while we need (C_4) for the entropy integral in the "standard case", the proof for the sieved variant goes similarly to the derivation of linear regression bounds (Corollary 6) and ignores (C_4) completely. Because of this, the results for the sieved case can be easily generalized to an unbounded domain (when $\boldsymbol{\mathcal{X}}$ is sub-Gaussian) by Theorem 1, but the standard case still needs a uniformly bounded hypothesis class to satisfy (C_4) without weakening the rate. This raises the question, whether sieved estimators have this benefit on the top of the rate improvement compared to the standard ones, or perhaps there might appear further improvements or better alternatives to the chaining technique on the "tail" (Theorem 10) in the future.

A Orlicz spaces of random vectors

In this appendix, we shortly review a few useful properties of the Ψ_q-Orlicz norm $\|\cdot\|_{\Psi_q}$. We start with the basic characteristics by Lemma 12.

Lemma 12. *Let $p \geq q \geq 1$, and $\boldsymbol{\mathcal{W}}, \boldsymbol{\mathcal{Z}} \in \mathbb{R}^d$ be random vectors such that $\|\boldsymbol{\mathcal{W}}\|_{\Psi_q} < \infty$ and $\|\boldsymbol{\mathcal{Z}}\|_{\Psi_p} < \infty$. Then the following statements hold:*

(a) $\|\boldsymbol{\mathcal{Z}}\|_{\Psi_q} \leq \|\boldsymbol{\mathcal{Z}}\|_{\Psi_p}$,

(b) $\|c\boldsymbol{\mathcal{W}}\|_{\Psi_q} = |c| \, \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}$, $c \in \mathbb{R}$, $\|\boldsymbol{\mathcal{W}} + \boldsymbol{\mathcal{Z}}\|_{\Psi_q} \leq \|\boldsymbol{\mathcal{W}}\|_{\Psi_q} + \|\boldsymbol{\mathcal{Z}}\|_{\Psi_q}$,

(c) $\mathbb{P}\big\{ \|\boldsymbol{\mathcal{W}}\| \geq t \big\} \leq 2e^{-t^q / \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^q}$ *for all* $t \geq 0$,

(d) $\mathbb{E}\big[\|\boldsymbol{\mathcal{W}}\|^s \big] \leq 2\big(s/(eq)\big)^{s/q} \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^s$ *for all* $s > 0$.

Proof. For (a) observe that $(x, q) \mapsto e^{x^q}$ is monotone increasing in q for any $x \geq 0$.

For the first claim of (b) simply use the definition of $\|\cdot\|_{\Psi_q}$. For the second claim, use (a) with $p \geq q$ and notice that Ψ_q is convex by $q \geq 1$.

For (c), use the Chernoff bound as

$$\mathbb{P}\{\|\boldsymbol{\mathcal{W}}\| > t\} \leq e^{-t^q / \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^q} \, \mathbb{E}\Big[e^{\|\boldsymbol{\mathcal{W}}\|^q / \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^q} \Big] \leq 2e^{-t^q / \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^q} .$$

For (d), use $x^z \leq (z/e)^z \, e^x$, $x \geq 0$, $z > 0$ (Buldygin and Kozachenko, 2000, in the proof of Lemma 1.4) by $x \doteq (\|\mathcal{W}\| \,/\, \|\mathcal{W}\|_{\Psi_q})^q$, $z \doteq s/q$, and take the expectation of both sides as

$$\mathbb{E}\big[\|\boldsymbol{\mathcal{W}}\|^s \big] \leq \big(s/(eq)\big)^{s/q} \, \mathbb{E}\Big[e^{\|\boldsymbol{\mathcal{W}}\|^q / \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^q} \Big] \, \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^s \leq 2\big(s/(eq)\big)^{s/q} \, \|\boldsymbol{\mathcal{W}}\|_{\Psi_q}^s . \qquad \blacksquare$$

Notice that by Lemma 12b, $\|\cdot\|_{\Psi_q}$ is indeed a norm for $q \geq 1$ on the space of random vectors with $\|\boldsymbol{\mathcal{W}}\|_{\Psi_q} < \infty$ (by also using $\|\boldsymbol{\mathcal{W}}\|_{\Psi_q} = 0$ if and only if $\boldsymbol{\mathcal{W}} = \mathbf{0}$).

Next, Lemma 13 provides a large deviation inequality for the sum of independent random variables with finite Ψ_q-Orlicz norm for the subexponential ($q = 1$) and sub-Gaussian ($q = 2$) cases.

Lemma 13. *Let $q \in \{1,2\}$, and $\boldsymbol{\mathcal{W}}_1, \dots, \boldsymbol{\mathcal{W}}_n \in \mathbb{R}^d$ be independent random variables with $\mathbb{E}[\boldsymbol{\mathcal{W}}_i] = 0$ and $\|\boldsymbol{\mathcal{W}}_i\|_{\Psi_q} < \infty$ for all $i = 1, \dots, n$. Then,*

$$\left\| \sum_{i=1}^{n} \boldsymbol{\mathcal{W}}_i \right\|_{\Psi_q} \leq 4\sqrt[q]{d} \left(\sum_{i=1}^{n} \|\boldsymbol{\mathcal{W}}_i\|_{\Psi_q} \right)^{1/2}.$$

Proof. For $d = 1$ and $q = 1$, see Corollary 3.5 of Buldygin and Kozachenko (2000). For $d = 1$ and $q = 2$, combine Lemmas 1.6 and 1.7 of Buldygin and Kozachenko (2000) with the remark at (2.4) of Boucheron et al. (2012). For $d > 1$ and $q \in \{1,2\}$, write $\boldsymbol{\mathcal{W}}_i = [\mathcal{W}_{i1} \dots \mathcal{W}_{id}]^\top \in \mathbb{R}^d$ for all $i = 1, \dots, n$, and notice that $\|\boldsymbol{\mathcal{W}}_i\|_q^q = \sum_{j=1}^d |\mathcal{W}_{ij}|^q$. Hence, by using $C \doteq 4\sqrt{\sum_{i=1}^n \|\boldsymbol{\mathcal{W}}_i\|_{\Psi_q}^2}$ and Hölder's inequality, we get

$$\mathbb{E}\left[e^{\left\|\sum_{i=1}^n \boldsymbol{\mathcal{W}}_i\right\|_q^q/(dC^q)}\right] = \mathbb{E}\left[e^{\sum_{j=1}^d \left|\sum_{i=1}^n \mathcal{W}_{ij}\right|^q/(dC^q)}\right] \leq \prod_{j=1}^{d} \mathbb{E}\left[e^{\left|\sum_{i=1}^n \mathcal{W}_{ij}\right|^q/C^q}\right]^{1/d} \leq 2\,,$$

where the last inequality follows from the $d = 1$ case. ■

To derive upper bounds for ERM estimators by Lemma 3, we apply Bernstein's inequality (Lemma 14) for the product of subexponential random variables $\mathcal{W}\mathcal{Z}$, which requires an "appropriate" bound on the higher moments as provided by Lemma 15. Here, "appropriate" means that the bound has to scale with the second moment of one multiplier, say $\mathcal{W}$, replacing $\|\mathcal{W}\|_{\Psi_q}^2$ by $\mathbb{E}[\mathcal{W}^2]$. The price we pay for this is only logarithmic in the kurtosis $\mathbb{K}_0[\mathcal{W}]$, which is crucial to our analysis for deriving excess risk upper bound for problems with highly-skewed distributions (see Section 1.2).

Lemma 14 (Bernstein's lemma)**.** *Consider a real valued random variable $\mathcal{W}$ satisfying $\mathbb{E}\big[|\mathcal{W}|^k\big] \leq (k!/2)v^2c^{k-2}$ for all $2 \leq k \in \mathbb{N}$. Then, for all $|s| < 1/c$,*

$$\ln \mathbb{E}\left[e^{s(\mathbb{E}[\mathcal{W}]-\mathcal{W})}\right] \leq \frac{s^2\, v^2}{2(1-|s|c)}\,.$$

Proof. See, for example Boucheron et al. (2012, Theorem 2.10) with $n \leftarrow 1$, and use $X_1 \leftarrow -\mathcal{W}$, $\lambda \leftarrow -s$ when $s < 0$. ■

Lemma 15. *Let $p, q \geq 1$ such that $\frac{1}{p} + \frac{1}{q} \leq 1$, and let $\mathcal{W}, \mathcal{Z} \in \mathbb{R}$ be two random variables such that $\mathbb{E}[\mathcal{W}^2] > 0$, $\|\mathcal{W}\|_{\Psi_q} \leq B$, and $\|\mathcal{Z}\|_{\Psi_p} \leq R$ with some $B, R > 0$. Then for all $2 \leq k \in \mathbb{N}$, we have*

$$\mathbb{E}\big[|\mathcal{W}\mathcal{Z}|^k\big] \leq (k!/2)\,\mathbb{E}\big[\mathcal{W}^2\big](2cR)^2\left(c^2BR\right)^{k-2},$$

where $c^{\min\{p,q\}} \doteq 2\ln\left(64\,\mathbb{K}_0[\mathcal{W}]\right)$.

Proof. Let $c > 0$ to be chosen later. Then by the Cauchy-Schwartz inequality, we have

$$\begin{aligned}
\mathbb{E}\big[|\mathcal{W}\mathcal{Z}|^k\big] &= \mathbb{E}\big[|\mathcal{W}\mathcal{Z}|^k \mathbb{I}\{|\mathcal{W}| \le cB\}\mathbb{I}\{|\mathcal{Z}| \le cR\}\big] \\
&\quad + \mathbb{E}\big[|\mathcal{W}\mathcal{Z}|^k \mathbb{I}\{|\mathcal{W}| \le cB\}\mathbb{I}\{|\mathcal{Z}| > cR\}\big] \\
&\quad + \mathbb{E}\big[|\mathcal{W}\mathcal{Z}|^k \mathbb{I}\{|\mathcal{W}| > cB\}\big] \\
&\le \mathbb{E}[\mathcal{W}^2](cR)^2(c^2BR)^{k-2} \\
&\quad + \mathbb{E}[\mathcal{W}^4]^{\frac{1}{2}}\,(cB)^{k-2}\,\mathbb{E}\big[\mathcal{Z}^{2k}\,\mathbb{I}\{|\mathcal{Z}| > cR\}\big]^{\frac{1}{2}} \\
&\quad + \mathbb{E}[\mathcal{W}^4]^{\frac{1}{2}}\,\mathbb{E}\big[\mathcal{W}^{4(k-2)}\mathcal{Z}^{4k}\big]^{\frac{1}{4}}\,\mathbb{P}\{|\mathcal{W}| > cB\}^{\frac{1}{4}} \\
&\le \mathbb{E}[\mathcal{W}^2]\Big((cR)^2(c^2BR)^{k-2} \\
&\quad + \mathbb{K}_0[\mathcal{W}]^{\frac{1}{2}}(cB)^{k-2}\,\mathbb{E}\big[\mathcal{Z}^{4k}\big]^{\frac{1}{4}}\,\mathbb{P}\{|\mathcal{Z}| > cR\}^{\frac{1}{4}} \\
&\quad + \mathbb{K}_0[\mathcal{W}]^{\frac{1}{2}}\,\mathbb{E}\big[\mathcal{W}^{8(k-2)}\big]^{\frac{1}{8}}\,\mathbb{E}\big[\mathcal{Z}^{8k}\big]^{\frac{1}{8}}\,\mathbb{P}\{|\mathcal{W}| > cB\}^{\frac{1}{4}}\Big)\,.
\end{aligned}$$

Now let $m \doteq \min\{p, q\}$, and set $c^m \doteq 4\ln\big(8\,\mathbb{K}_0^{1/2}[\mathcal{W}]\big)$, which satisfies $c^m \ge 8$ as $\mathbb{K}_0[\mathcal{W}] \ge 1$ by Jensen's inequality. Then apply Lemma 12c, $2(k/e)^k \le k!$, and Lemma 12d, to get

$$\begin{aligned}
\mathbb{E}\big[|\mathcal{W}\mathcal{Z}|^k\big] &\le \mathbb{E}[\mathcal{W}^2](cR)^2(c^2BR)^{k-2}\cdot \\
&\quad \left(1 + 2^{5/4}\,\mathbb{K}_0[\mathcal{W}]^{\frac{1}{2}}\Big(\big(\frac{4k}{ep\,c^p}\big)^{\frac{k}{p}}\,e^{-c^p/4} + \big(\frac{8(k-2)}{eq\,c^q}\big)^{\frac{k-2}{q}}\big(\frac{8k}{ep\,c^p}\big)^{\frac{k}{p}}\,e^{-c^q/4}\Big)\right) \\
&\le \mathbb{E}[\mathcal{W}^2](cR)^2(c^2BR)^{k-2}\left(1 + 2^{5/4}\,\mathbb{K}_0[\mathcal{W}]^{\frac{1}{2}}\,e^{-c^m/4}\Big((k/e)^k + \big(k/e\big)^{k(\frac{1}{q}+\frac{1}{p})}\Big)\right) \\
&\le \mathbb{E}[\mathcal{W}^2](cR)^2(c^2BR)^{k-2}\left(1 + 2^{5/4}\,\mathbb{K}_0[\mathcal{W}]^{\frac{1}{2}}\,e^{-c^m/4}\,k!\right) \\
&\le 2k!\,\mathbb{E}[\mathcal{W}^2](cR)^2\,(c^2BR)^{k-2}\,.
\end{aligned}$$

■

B Auxiliary tools

In this appendix we provide a few auxiliary results which are used for our proofs.

Lemma 16. *Let $p \in \mathbb{N} \cup \{\infty\}$ and $\mathcal{S} \subset \mathbb{R}^d$ with a finite radius under $\|\cdot\|_p$, that is suppose there exists $\boldsymbol{x}_* \in \mathbb{R}^d$ such that $\mathcal{S} \subseteq \{\boldsymbol{x} \in \mathbb{R}^d : \|\boldsymbol{x} - \boldsymbol{x}_*\|_p \le R\}$ for some $R > 0$. Then $\mathcal{H}_{\|\cdot\|_p}(\epsilon, \mathcal{S}) \le d\ln(3R/\epsilon)$ for all $\epsilon \in (0, 3R]$.*

Proof. By using the volume argument (e.g., see Pollard, 1990, proof of Lemma 4.1). ■

Lemma 17. *For $A \in \mathbb{R}^{n\times d}$, $\boldsymbol{b} \in \mathbb{R}^n$ and $r > 0$, $\big\|(rI_d + A^\top A)^{-1}A^\top\boldsymbol{b}\big\| \le \frac{\|\boldsymbol{b}\|}{2\sqrt{r}}$.*

Proof. Consider a thin singular value decomposition of A given as $A = USV^\top$, where $U \in \mathbb{R}^{n\times d}$ is semi-orthogonal ($U^\top U = I_d$), $S \in \mathbb{R}^{d\times d}$ is diagonal with nonegative elements, and $V \in \mathbb{R}^{d\times d}$ is orthogonal ($V^\top V = VV^\top = I_d$). Then

$$\Big\|(rI_d + A^\top A)^{-1}A^\top\boldsymbol{b}\Big\| = \Big\|V(rI_d + S^2)^{-1}SU^\top\boldsymbol{b}\Big\| \le \big\|(rI_d + S^2)^{-1}S\big\|\,\|\boldsymbol{b}\|\,,$$

where we used $\|U\| = \|V\| = 1$ (as $U^\top U = I_d$ implies $\|U\boldsymbol{x}\|^2 = \|\boldsymbol{x}\|^2$). Finally notice that

$$\left\|(rI_d + S^2)^{-1}S\right\| \le \max_{s\ge 0} \frac{s}{r+s^2} = \frac{1}{2\sqrt{r}}\,,$$

which proves the claim. ∎

Proof of Lemma 4 for $\boldsymbol{a}_n$. By the ERM property (1) of f_n and because the constant function $\boldsymbol{x} \mapsto \overline{\mathcal{Y}}$ is in $\mathcal{F}_{\text{aff}}^{L,n}$, we obtain $R_n(f_n) + \beta(f_n) \le R_n(\boldsymbol{x} \mapsto \overline{\mathcal{Y}}) + \beta(\boldsymbol{x} \mapsto \overline{\mathcal{Y}}) + \alpha$. This can be rearranged into $\frac{1}{n}\sum_{i=1}^n |\boldsymbol{a}_n^\top(\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})|^2 \le \frac{2}{n}\sum_{i=1}^n \boldsymbol{a}_n^\top(\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})(\mathcal{Y}_i - \overline{\mathcal{Y}}) + \alpha$, where we also used $\beta(\cdot) \ge 0$ and $\beta(\boldsymbol{x} \mapsto \overline{\mathcal{Y}}) = 0$.

Using this, $\hat{\Sigma} \doteq \frac{1}{n}\sum_{i=1}^n (\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})(\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})^\top$, and $2ab \le \frac{a^2}{2} + 2b^2$, we obtain

$$\boldsymbol{a}_n^\top \hat{\Sigma} \boldsymbol{a}_n = \frac{1}{n}\sum_{i=1}^n |\boldsymbol{a}_n^\top(\boldsymbol{\mathcal{X}}_i - \overline{\boldsymbol{\mathcal{X}}})|^2 \le \frac{4}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2 + \alpha \le \frac{5}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2\,,$$

which can be transformed to $\boldsymbol{a}_n^\top \Sigma \boldsymbol{a}_n \le \frac{1}{r}\big(\frac{5}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2 + \boldsymbol{a}_n^\top(r\Sigma - \hat{\Sigma})\boldsymbol{a}_n\big)$ with any $r \in (0,1]$. In particular, with $r \doteq 1/2$, we have

$$\boldsymbol{a}_n^\top \Sigma \boldsymbol{a}_n \le \frac{10}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2 + \mathcal{W}\,, \quad \mathcal{W} \doteq \boldsymbol{a}_n^\top\big((1/2)\Sigma - \hat{\Sigma}\big)\boldsymbol{a}_n\,. \tag{16}$$

Then, observe that $\frac{1}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2 = \frac{1}{n}\sum_{i=1}^n |\mathcal{Y}_i - \mathbb{E}\mathcal{Y}|^2 - |\overline{\mathcal{Y}} - \mathbb{E}\mathcal{Y}|^2$, so by Lemma 12c, we get with probability at least $1 - \gamma/3$ that $\frac{10}{n}\sum_{i=1}^n |\mathcal{Y}_i - \overline{\mathcal{Y}}|^2 \le 10\sigma^2 \ln(6/\gamma)$.

Next, notice that

$$\begin{aligned}\mathcal{W} &= (1/2)\mathbb{E}\big[|\boldsymbol{a}_n^\top(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})|^2\big] - \frac{1}{n}\sum_{i=1}^n |\boldsymbol{a}_n^\top(\boldsymbol{\mathcal{X}}_i - \mathbb{E}\boldsymbol{\mathcal{X}})|^2 + |\boldsymbol{a}_n^\top(\overline{\boldsymbol{\mathcal{X}}} - \mathbb{E}\boldsymbol{\mathcal{X}})|^2 \\ &\le \sup_{\boldsymbol{a}:\|\boldsymbol{a}\|\le L} \Gamma(\boldsymbol{a}) + L^2\|\overline{\boldsymbol{\mathcal{X}}} - \mathbb{E}\boldsymbol{\mathcal{X}}\|^2\,,\end{aligned}$$

where $\Gamma(\boldsymbol{a}) \doteq (1/2)\mathbb{E}\big[|\boldsymbol{a}^\top(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})|^2\big] - \frac{1}{n}\sum_{i=1}^n |\boldsymbol{a}^\top(\boldsymbol{\mathcal{X}}_i - \mathbb{E}\boldsymbol{\mathcal{X}})|^2$. Then, by Lemma 13, we have $\|\overline{\boldsymbol{\mathcal{X}}} - \mathbb{E}\boldsymbol{\mathcal{X}}\|_{\Psi_2} \le 4\rho\sqrt{d/n}$, so Lemma 12c also implies with probability at least $1 - \gamma/3$ that $L^2\|\overline{\boldsymbol{\mathcal{X}}} - \mathbb{E}\boldsymbol{\mathcal{X}}\|^2 \le \frac{16d(L\rho)^2}{n}\ln(6/\gamma)$.

It remains to bound $\sup_{\boldsymbol{a}:\|\boldsymbol{a}\|\le L}\Gamma(\boldsymbol{a})$ with probability at least $1 - \gamma/3$, for which we use *Lemma 3* and Theorem 11. First, use $\mathcal{Y} \leftarrow 0$, $f_* \leftarrow (\boldsymbol{x} \mapsto 0)$, $\Delta_{f,f_*} \leftarrow \boldsymbol{a}^\top(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})$ to satisfy the requirements of Lemma 3 with $\|\Delta_{f,f_*}\|_{\Psi_2} \le L\rho$, $r_0 \to 0$, $C = 1$, $R = L\rho$ due to $\mathcal{Z}(f, f_*) = \Delta_{f,f_*}^2$. Hence, we get $\sup_{\boldsymbol{a}:\|\boldsymbol{a}\|\le L}\mathbb{E}\big[e^{\mathbb{E}[|\boldsymbol{a}^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}})|^2]/(2\theta) - |\boldsymbol{a}^\top(\boldsymbol{\mathcal{X}}-\mathbb{E}\boldsymbol{\mathcal{X}})|^2/\theta}\big] \le 1$ with $\theta = 72(L\rho)^2\ln(23/\eta_\mu)$ by $t \leftarrow 9$ and $\sup_{\boldsymbol{a}\neq\boldsymbol{0}}\mathbb{K}_0[\boldsymbol{a}^\top(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})] \le 2/\eta_\mu^2$ due to Lemma 12d with $s = 4$ and $q = 2$. This further implies $\sup_{\boldsymbol{a}:\|\boldsymbol{a}\|\le L}\mathbb{E}[e^{\Gamma(\boldsymbol{a})/(\theta/n)}] \le 1$ due to the i.i.d. property of $\boldsymbol{\mathcal{X}}, \boldsymbol{\mathcal{X}}_1, \ldots, \boldsymbol{\mathcal{X}}_n$, and gives us condition Theorem 11c for all $\boldsymbol{a}$ having $\|\boldsymbol{a}\| \le L$.

Furthermore, for any $\boldsymbol{a}, \hat{\boldsymbol{a}}$ with $\|\boldsymbol{a}\| \le L$ and $\|\hat{\boldsymbol{a}}\| \le L$, and $\Lambda(\boldsymbol{a}) \doteq \Gamma(\boldsymbol{a}) - \mathbb{E}[\Gamma(\boldsymbol{a})]$, we get by using $a^2 - b^2 = (a-b)(a+b)$ that

$$\begin{aligned}&\Lambda(\boldsymbol{a}) - \Lambda(\hat{\boldsymbol{a}}) \\ &= \mathbb{E}\Big[|\boldsymbol{a}^\top(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})|^2 - |\hat{\boldsymbol{a}}^\top(\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}})|^2\Big] + \frac{1}{n}\sum_{i=1}^n |\hat{\boldsymbol{a}}^\top(\boldsymbol{\mathcal{X}}_i - \mathbb{E}\boldsymbol{\mathcal{X}})|^2 - |\boldsymbol{a}^\top(\boldsymbol{\mathcal{X}}_i - \mathbb{E}\boldsymbol{\mathcal{X}})|^2 \\ &\le \frac{\|\boldsymbol{a} - \hat{\boldsymbol{a}}\|}{L}(2L^2)\Big(\mathbb{E}\big[\,\|\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}}\|^2\,\big] + \frac{1}{n}\sum_{i=1}^n \|\boldsymbol{\mathcal{X}}_i - \mathbb{E}\boldsymbol{\mathcal{X}}\|^2\Big) = \psi(\boldsymbol{a}, \hat{\boldsymbol{a}})\,\tau(\mathcal{D}_n)\,,\end{aligned}$$

with $\psi(\boldsymbol{a}, \hat{\boldsymbol{a}}) \doteq \|\boldsymbol{a} - \hat{\boldsymbol{a}}\|/L$ and $\tau(\mathcal{D}_n) \doteq 2L^2\big(\mathbb{E}\big[\,\|\boldsymbol{\mathcal{X}} - \mathbb{E}\boldsymbol{\mathcal{X}}\|^2\,\big] + \frac{1}{n}\sum_{i=1}^n \|\boldsymbol{\mathcal{X}}_i - \mathbb{E}\boldsymbol{\mathcal{X}}\|^2\,\big)$. By Lemma 12c with $\|\sqrt{\tau(\mathcal{D}_n)}\|_{\Psi_2} \le 2L\rho$, we have $\tau(\mathcal{D}_n) \le 4(L\rho)^2 \ln(12/\gamma) \le 6(L\rho)^2 \ln(6/\gamma)$ with probability at least $1-\gamma/6$. Hence, we also have condition Theorem 11a. Furthermore, by Lemma 16, using $\sup_{\boldsymbol{a}:\|\boldsymbol{a}\|\le L} \psi(\boldsymbol{a}, \boldsymbol{0}) \le 1$, we also have $\mathcal{H}_\psi(\epsilon, \{\boldsymbol{a} : \|\boldsymbol{a}\| \le L\}) \le d\ln(3/\epsilon)$ for all $\epsilon \in (0, 3]$.

Then, by using Theorem 11 with $\epsilon \doteq \delta \doteq 1/n \in (0, 3]$ and $\gamma \leftarrow \gamma/3$, we get with probability at least $1 - \gamma/3$ that

$$\begin{aligned}\sup_{\boldsymbol{a}:\|\boldsymbol{a}\|\le L} \Gamma(\boldsymbol{a}) &\le \Big(\frac{\theta\mathcal{H}_\psi(\epsilon, \{\boldsymbol{a} : \|\boldsymbol{a}\| \le L\})}{n} + 48\epsilon(L\rho)^2\Big)\ln(6/\gamma) \\ &\le \frac{(L\rho)^2}{n}\Big(72d\ln(23/\eta_\mu)\ln(3n) + 48\Big)\ln(6/\gamma)\,.\end{aligned}$$

Finally, combining the three probabilistic bounds with (16), we get with probability at least $1 - \gamma$ that

$$\rho^2\eta_\mu \|\boldsymbol{a}_n\|^2 \le \boldsymbol{a}_n^\top \Sigma \boldsymbol{a}_n \le \Big(10\sigma^2 + \frac{8(L\rho)^2}{n}\big(11d\ln(23/\eta_\mu)\ln(3n) + 6\big)\Big)\ln(6/\gamma)\,,$$

which completes the proof after rearrangement. ■

References

G. Balázs, A. György, and C. Szepesvári. Near-optimal max-affine estimators for convex regression. In G. Lebanon and S. Vishwanathan, editors, *The 18th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 38 of *JMLR W&CP*, pages 56–64, 2015.

P. L. Bartlett and S. Mendelson. Empirical minimization. *Probability Theory and Related Fields*, 135(3):311–334, 2006.

P. L. Bartlett, O. Bousquet, and S. Mendelson. Local rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.

P. L. Bartlett, M. I. Jordan, and J. D. McAuliffe. Convexity, classification, and risk bounds. *Journal of the American Statistical Association*, 101(473):138–156, 2006.

S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A nonasymptotic theory of independence*. Clarendon Press, 2012.

V. V. Buldygin and Y. V. Kozachenko. *Metric characterization of random variables and random processes*, volume 188 of *Translations of Mathematical Monographs*. American Mathematical Society, Providence, RI, 2000.

N. Cesa-Bianchi and G. Lugosi. Minimax regret under log loss for general classes of experts. In S. Ben-David and P. M. Long, editors, *The 12th Conference on Computational Learning Theory (COLT)*, pages 12–18, 1999.

R. M. Dudley. *Uniform Central Limit Theorems*. Cambridge University Press, 1999.

P. D. Grünwald and N. A. Mehta. Fast rates with unbounded losses. *CoRR*, 2016. http://arxiv.org/abs/1605.00252v1.

L. Györfi and M. Wegkamp. Quantization for nonparametric regression. *IEEE Transactions on Information Theory*, 54(2):867–874, 2008.

L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer-Verlag, 2002.

A. E. Hoerl and R. W. Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.

D. Hsu, S. M. Kakade, and T. Zhang. Random design analysis of ridge regression. *Foundations of Computational Mathematics*, 14:569–600, 2014.

R. Huang and C. Szepesvári. A finite-sample generalization bound for semiparametric regression: Partially linear models. In S. Kaski and J. Corander, editors, *The 17th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 33 of *JMLR W&CP*, pages 402–410, 2014.

V. Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*. Springer, 2011.

G. Lecué and S. Mendelson. Learning subgaussian classes: Upper and minimax bounds. *CoRR*, 2013. http://arxiv.org/abs/1305.4825v1.

T. Liang, A. Rakhlin, and K. Sridharan. Learning with squared loss: Localization through offset rademacher complexity. In P. Grünwald, E. Hazan, and S. Kale, editors, *The 28th Conference on Learning Theory (COLT)*, volume 40 of *JMLR W&CP*, 2015.

E. Lim. On convergence rates of convex regression in multiple dimensions. *INFORMS Journal of Computing*, 26(3):616–628, 2014.

N. A. Mehta and R. C. Williamson. From stochastic mixability to fast rates. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 27 (NIPS)*, pages 1197–1205, 2014.

S. Mendelson. Learning without concentration. In M. F. Balcan, V. Feldman, and C. Szepesvári, editors, *The 27th Conference on Learning Theory (COLT)*, volume 35 of *JMLR W&CP*, pages 25–39, 2014.

D. Pollard. *Empirical Processes: Theory and Applications*. Institute of Mathematical Statistics, 1990.

O. Shamir. The sample complexity of learning predictors with the squared loss. *Journal of Machine Learning Research (JMLR)*, 16:3475–3486, 2015.

S. van de Geer. *Empirical Processes in M-Estimation*. Cambridge University Press, 2000.

A. van der Vaart and J. A. Wellner. A local maximal inequality under uniform entropy. *Electronic Journal of Statistics*, 5:192–203, 2011.